arXiv:1507.00623v1 [cs.LO] 2 Jul 2015

Assume-Admissible Synthesis*

Romain Brenguier, Jean-François Raskin, Ocan Sankur

Université Libre de Bruxelles, Brussels, Belgium

July 3, 2015

Abstract

In this paper, we introduce a novel rule for synthesis of reactive systems, applicable to systems made of n components which have each their own objectives. It is based on the notion of *admissible* strategies. We compare our novel rule with previous rules defined in the literature, and we show that contrary to the previous proposals, our rule defines sets of solutions which are *rectangular*. This property leads to solutions which are robust and resilient. We provide algorithms with optimal complexity and also an abstraction framework.

1 Introduction

The automatic synthesis of reactive systems has recently attracted a considerable attention. The theoretical foundations of most of the contributions in this area rely on two-player zero sum games played on graphs: one player (player 1) models the system to synthesize, and the other player (player 2) models its environment. The game is zero-sum: the objective of player 1 is to enforce the specification of the system while the objective of player 2 is the *negation* of this specification. This is a *worst-case* assumption: because the cooperation of the environment cannot be assumed, we postulate that it is *antagonistic*.

A fully adversarial environment is usually a bold abstraction of reality. Nevertheless, it is popular because it is *simple* and *sound*: a *winning strategy* against an antagonistic player is winning against *any* environment which pursues its own objective. But this approach may fail to find a winning strategy even if there exist solutions when the objective of the environment is taken into account. Also, this model is for two players only: system *vs* environment. In practice, both the system and the environment may be composed of several parts to be constructed individually or whose objectives should be considered one at a time. It is thus crucial to take into account different players' objectives when synthesizing strategies; accordingly, alternative notions have been proposed in the literature.

A first classical alternative is to *weaken* the winning condition of player 1 using the objective of the environment, requiring the system to win only when the environment meets its objective. This approach together with its weaknesses have been discussed in [3], we will add to that later in the paper. A second alternative is to use concepts from n-players *non*-zero sum games. This is the approach taken both by *assume-guarantee synthesis* [6] (AG), and by *rational synthesis* [16] (RS). AG relies on *secure equilibria* [8] (SE), a refinement of Nash equilibria [25] (NE). In SE, objectives are lexicographic: players first try to maximize their own specifications, and then try to falsify the specifications of others. It is shown in [8] that SE are those NE which represent enforceable contracts between the two players. In RS, the system is assumed to be monolithic and the environment is made of components that are *partially controllable*. In RS, we search for a profile of strategies where the system ensures its objective and the players that model the environment are given an "*acceptable*" strategy profiles, from which it is assumed that they will not deviate. "Acceptable" is formalized either by NE, dominating strategies (Dom), or subgame perfect equilibria (SPE).

*Supported by the ERC starting grant inVEST (FP7-279499)

Contributions. As a first and central contribution, we propose a novel notion of synthesis where we take into account different players' objectives using the concept of *admissible* strategies [1, 2, 4]. For a player with objective ϕ, a strategy σ is *dominated* by σ' if σ' does as well as σ w.r.t. ϕ against all strategies of the other players, and better for some of those strategies. A strategy σ is *admissible* if it is *not* dominated by another strategy. In [2], the admissibility notion was lifted to games played on graphs, and algorithmic questions left open were solved in [4], with the goal of *model checking* the set of runs that survive the iterative elimination of dominated strategies. Here, we use this notion to derive a meaningful notion to *synthesize* systems with several players, with the following idea. *Rational* players should only play admissible strategies since dominated strategies are clearly *suboptimal*. In *assume-admissible synthesis* (AA), we make the assumption that players play admissible strategies. Then for each player, we search for an admissible strategy that is *winning* against all admissible strategies of other players. AA is *sound*: any strategy profile that is winning against admissible strategies of other players, satisfies the objectives of all the players (Theorem 1).

As a second contribution, we compare the different synthesis rules. First we apply all the rules on a simple but representative example, and show the main advantages of AA w.r.t. the other rules. Then we compare systematically the different approaches. We show when a solution for one rule implies a solution for another rule and we prove that, contrary to other rules, AA yields rectangular sets of solutions (Theorem 3). We argue that the rectangularity property is essential for practical applications.

As a third contribution, we provide algorithms to decide the existence of assume-admissible winning strategy profiles and prove the optimal complexity of our algorithm (Theorem 4): PSPACE-complete for Müller, and PTIME for Büchi objectives.

As a last important contribution, we provide an abstraction framework which allows us to define sufficient conditions to compute sets of winning assume-admissible strategies for each player in the game compositionally (Theorem 5).

Additional pointers to related works. We have already mentioned assume-guarantee synthesis [6] and rational synthesis [16, 21]. Those are the closest related works to ours as they pursue the same scientific objective: to synthesis strategy profiles for non-zero sum multi-player games by taking into account the specification of each player. As those works are defined for similar formal setting, we are able to provide formal statements in the core of the paper that add elements of comparison with our work.

In [15], Faella studies several alternatives to the notion of winning strategy including the notion of admissible strategy. His work is for two-players only, and only the objective of one player is taken into account, the objective of the other player is left unspecified. Faella uses the notion of admissibility to define a notion of *best-effort* in synthesis while we use the notion of admissibility to take into account the objectives of the other players in an n player setting where each player has his own objective.

The notion of admissible strategy is definable in strategy logics [9, 24] and decision problems related to the AA rule can be reduced to satisfiability queries in such logics. Nevertheless this would not lead to worst-case optimal algorithms. Based on our previous work [4], we develop in this paper worst-case optimal algorithms.

In [12], Damm and Finkbeiner use the notion of *dominant strategy* to provide a compositional semi-algorithm for the (undecidable) distributed synthesis problem. So while we use the notion of admissible strategy, they use a notion of dominant strategy. The notion of dominant strategy is *strictly stronger*: every dominant strategy is admissible but an admissible strategy is not necessary dominant. Also, in multiplayer games with omega-regular objectives with complete information (as considered here), admissible strategies are always guaranteed to exist [2] while it is not the case for dominant strategies. We will show in an example that the notion of dominant strategy is too strong for our purpose. Also, note that the objective of Damm and Finkbeiner is different from ours: they use dominance as a mean to formalize a notion of *best-effort* for components of a distributed system w.r.t. their common objective, while we use admissibility to take into account the objectives of the other components when looking for a winning strategy for one component to enforce its own objective. Additionally, our formal setting is different from their setting in several respects. First, they consider zero-sum games between a distributed team of players (processes) against a unique environment, each player in the team has the same specification (the specification of the distributed system to synthesize) while the environment is considered as adversarial and so its specification is the negation of

the specification of the system. In our case, each player has his *own* objective and we do not distinguish between protagonist and antagonist players. Second, they consider distributed synthesis: each individual process has its own view of the system while we consider games with perfect information in which all players have a complete view of the system state. Finally, let us point out that Damm and Finkbeiner use the term *admissible* for specifications and *not* for strategies (as already said, they indeed consider dominant strategies and not admissible strategies). In our case, we use the notion of *admissible* strategy which is classical in game theory, see e.g. [17, 1]. This vocabulary mismatch is unfortunate but we decided to stick to the term of "admissible strategy" which is well accepted in the literature, and already used in several previous works on (multi-player) games played on graphs [2, 15, 4].

Structure of the paper. Sect. 2 contains definitions. In Sect. 3, we review synthesis rules introduced in the literature and define assume-admissible synthesis. In Sect. 4, we consider an example; this allows us to underline some weaknesses of the previous rules. Sect. 5 presents a formal comparison of the different rules. Sect. 6 contains algorithms for Büchi and Müller objectives, and Sect. 7 abstraction techniques applied to our rule.

2 Definitions

A *turn-based multiplayer arena* is a tuple $\mathsf{A} = \langle \mathcal{P}, (\mathsf{S}_i)_{i\in\mathcal{P}}, s_{\mathsf{init}}, (\mathsf{Act}_i)_{i\in\mathcal{P}}, \delta\rangle$ where $\mathcal{P}$ is a finite set of players; for $i \in \mathcal{P}$, S_i is a finite set of player-i states; we let $\mathsf{S} = \biguplus_{i\in\mathcal{P}} \mathsf{S}_i$; $s_{\mathsf{init}} \in \mathsf{S}$ is the initial state; for every $i \in \mathcal{P}$, Act_i is the set of player-i actions; we let $\mathsf{Act} = \bigcup_{i\in\mathcal{P}} \mathsf{Act}_i$; and $\delta\colon \mathsf{S} \times \mathsf{Act} \mapsto \mathsf{S}$ is the transition function. A *run* ρ is a sequence of alternating states and actions $\rho = s_1a_1s_2a_2\ldots \in (\mathsf{S}\cdot\mathsf{Act})^\omega$ such that for all $i \geq 1$, $\delta(s_i, a_i) = s_{i+1}$. We write $\rho_i = s_i$, and $\mathsf{act}_i(\rho) = a_i$. A *history* is a finite prefix of a run ending in a state. We denote by $\rho_{\leq k}$ the history $s_1a_1\ldots s_k$; and write $\mathsf{last}(\rho_{\leq k}) = s_k$, the last state of the history. The set of states *occurring infinitely often* in a run ρ is $\mathrm{Inf}(\rho) = \{s \in \mathsf{S} \mid \forall j \in \mathbb{N}.\ \exists i > j, \rho_i = s\}$.

A *strategy* of player i is a function $\sigma_i : (\mathsf{S}^*\cdot\mathsf{S}_i) \to \mathsf{Act}_i$. A *strategy profile* for the set of players $P \subseteq \mathcal{P}$ is a tuple of strategies, one for each player of P. We write $-i$ for the set $\mathcal{P}\setminus\{i\}$. Let $\Sigma_i(\mathsf{A})$ be the set of the strategies of player i in A, written Σ_i if A is clear from context, and Σ_P the strategy profiles of $P \subseteq \mathcal{P}$.

A run ρ is *compatible* with strategy σ for player i if for all $j \geq 1$, $\rho_j \in \mathsf{S}_i$ implies $\mathsf{act}_j(\rho) = \sigma(\rho_{\leq j})$. It is compatible with strategy profile $\sigma_{\mathcal{P}}$ if it is compatible with each σ_i for $i \in \mathcal{P}$. The *outcome* of a strategy profile $\sigma_{\mathcal{P}}$ is the unique run compatible with $\sigma_{\mathcal{P}}$ starting at s_{init}, denoted $\mathsf{Out}_{\mathsf{A}}(\sigma_{\mathcal{P}})$. We write $\mathsf{Out}_{\mathsf{A},s}(\sigma_{\mathcal{P}})$ for the outcome starting at state s. Given $\sigma_P \in \Sigma_P$ with $P \subseteq \mathcal{P}$, let $\mathsf{Out}_{\mathsf{A}}(\sigma_P)$ denote the set of runs compatible with σ_P, and extend it to $\mathsf{Out}_{\mathsf{A}}(\Sigma')$ where Σ' is a set of strategy profiles. For $E \subseteq \mathsf{S}_i \times \mathsf{Act}_i$, let $\mathsf{Strat}_i(E)$ denote the set of player-i strategies σ whose compatible outcomes use action a from a state s only if $(s, a) \in E$.

An *objective* ϕ is a subset of runs. A strategy σ_i of player i is *winning* for objective ϕ_i if for all $\sigma_{-i} \in \Sigma_{-i}$, $\mathsf{Out}_{\mathsf{A}}(\sigma_i, \sigma_{-i}) \in \phi_i$. A *game* is an arena equipped with an objective for each player, written $\mathsf{G} = \langle \mathsf{A}, (\phi_i)_{i\in\mathcal{P}}\rangle$ where for each player i, ϕ_i is an objective. Given a strategy profile σ_P for the set of players P, we write $\mathsf{G}, \sigma_P \models \phi$ if $\mathsf{Out}_{\mathsf{A}}(\sigma_P) \subseteq \phi$. We write $\mathsf{Out}_{\mathsf{G}}(\sigma_P) = \mathsf{Out}_{\mathsf{A}}(\sigma_P)$, and $\mathsf{Out}_{\mathsf{G}} = \mathsf{Out}_{\mathsf{G}}(\Sigma)$ for $\Sigma \subseteq \Sigma_i$. For any coalition $C \subseteq \mathcal{P}$, and objective ϕ, we denote by $\mathsf{Win}_C(\mathsf{A}, \phi)$ the set of states s such that there exists $\sigma_C \in \Sigma_C$ with $\mathsf{Out}_{\mathsf{G},s}(\sigma_C) \subseteq \phi$.

Although we prove some of our results for general objectives, we give algorithms for ω-regular objectives represented by Muller conditions. A Muller condition is given by a family $\mathcal{F}$ of sets of states: $\phi_i = \{\rho \mid \mathrm{Inf}(\rho) \in \mathcal{F}\}$. Following [19], we assume that $\mathcal{F}$ is given by a Boolean circuit whose inputs are S, which evaluates to true exactly on valuations encoding subsets $S \in \mathcal{F}$. We also use linear temporal logic (LTL) [27] to describe objectives. LTL formulas are defined by $\phi := \mathtt{G}\phi \mid \mathtt{F}\phi \mid \mathtt{X}\phi \mid \phi\mathtt{U}\phi \mid \phi\mathtt{W}\phi \mid S$ where $S \subseteq \mathsf{S}$ (We refer to [14] for the semantics.) We consider the special case of Büchi objectives, given by $\mathtt{GF}(B) = \{\rho \mid B \cap \mathrm{Inf}(\rho) \neq \varnothing\}$. Boolean combinations of formulas $\mathtt{GF}(S)$ define Muller conditions representable by polynomial-size circuits.

In any game G, a player i strategy σ_i is *dominated* by σ'_i if for all $\sigma_{-i} \in \Sigma_{-i}$, $\mathsf{G}, \sigma_i, \sigma_{-i} \models \phi_i$ implies $\mathsf{G}, \sigma'_i, \sigma_{-i} \models \phi_i$ and there exists $\sigma_{-i} \in \Sigma_{-i}$, such that $\mathsf{G}, \sigma'_i, \sigma_{-i} \models \phi_i$ and $\mathsf{G}, \sigma_i, \sigma_{-i} \not\models \phi_i$, (this is classically called *weak* dominance, but we call it dominance for simplicity). A strategy which is not dominated is

admissible. Thus, admissible strategies are maximal, and incomparable, with respect to the dominance relation. We write $\mathsf{Adm}_i(\mathsf{G})$ for the set of *admissible* strategies in Σ_i, and $\mathsf{Adm}_P(\mathsf{G}) = \prod_{i \in P} \mathsf{Adm}_i(G)$ the product of the sets of admissible strategies for $P \subseteq \mathcal{P}$.

Strategy σ_i is *dominant* (Dom) if for all σ_i' and σ_{-i}, $\mathsf{G}, \sigma_i', \sigma_{-i} \models \phi_i$ implies $\mathsf{G}, \sigma_i, \sigma_{-i} \models \phi_i$. The set of dominant strategies for player i is written $\mathsf{Dom}_i(\mathsf{G})$. A *Nash equilibrium* (NE) for G is a strategy profile $\sigma_{\mathcal{P}}$ such that for all $i \in \mathcal{P}$, and $\sigma_i' \in \Sigma_i$, $\mathsf{G}, \sigma_{-i}, \sigma_i' \models \phi_i$ implies $\mathsf{G}, \sigma_{\mathcal{P}} \models \phi_i$; thus no player can improve its outcome by deviating from the prescribed strategy. A Nash equilibrium for G from s, is a Nash equilibrium for G where the initial state is replaced by s. A *subgame-perfect equilibrium* (SPE) for G is a strategy profile $\sigma_{\mathcal{P}}$ such that for all histories h, $(\sigma_i \circ h)_{i \in \mathcal{P}}$ is a Nash equilibrium in G from state $\mathsf{last}(h)$, where given a strategy σ, $\sigma \circ h$ denotes the strategy $\mathsf{last}(h) \cdot h' \mapsto \sigma(h \cdot h')$.

3 Synthesis Rules

In this section, we review synthesis rules proposed in the literature, and introduce a novel one: the *assume-admissible* synthesis rule (AA). Unless stated otherwise, we fix for this section a game G, with players $\mathcal{P} = \{1, \ldots, n\}$ and their objectives $\phi_1, \ldots, \phi_n$.

Rule Coop: The objectives are *achieved cooperatively* if there is a strategy profile $\sigma_{\mathcal{P}} = (\sigma_1, \sigma_2, \ldots, \sigma_n)$ such that $\mathsf{G}, \sigma_{\mathcal{P}} \models \bigwedge_{i \in \mathcal{P}} \phi_i$.

This rule [23, 10] asks for a strategy profile that *jointly* satisfies the objectives of all the players. This rule makes *very strong assumptions*: players fully cooperate and strictly follow their respective strategies. This concept is *not robust* against deviations and postulates that the behavior of every component in the system is *controllable*. This weakness is well-known: see e.g. [6] where the rule is called *weak co-synthesis*.

Rule Win. The objectives are *achieved adversarially* if there is a strategy profile $\sigma_{\mathcal{P}} = (\sigma_1, \ldots, \sigma_n)$ such that for all $i \in \mathcal{P}$, $\mathsf{G}, \sigma_i \models \phi_i$.

This rule does *not* require any cooperation among players: the rule asks to synthesize for each player i a strategy which enforces his/her objective ϕ_i against all possible strategies of the other players. Strategy profiles obtained by Win are extremely *robust*: each player is able to ensure his/her objective no matter how the other players behave. Unfortunately, this rule is often not applicable in practice: often, none of the players has a winning strategy against *all* possible strategies of the other players. The next rules soften this requirement by taking into account the objectives of other players.

Rule Win-under-Hyp: Given a two-player game G with $\mathcal{P} = \{1, 2\}$ in which player 1 has objective ϕ_1, player 2 has objective ϕ_2, player 1 can *achieve adversarially* ϕ_1 under *hypothesis* ϕ_2, if there is a strategy σ_1 for player 1 such that $\mathsf{G}, \sigma_1 \models \phi_2 \rightarrow \phi_1$.

The rule *winning under hypothesis* applies for two-player games only. Here, we consider the synthesis of a strategy for player 1 against player 2 under the hypothesis that player 2 behaves according to his/her specification. This rule is a relaxation of the rule Win as player 1 is *only* expected to win when player 2 plays so that the outcome of the game satisfies ϕ_2. While this rule is often reasonable, it is *fundamentally* plagued by the following problem: instead of trying to satisfy ϕ_1, player 1 could try to falsify ϕ_2, see e.g. [3]. This problem disappears if player 2 has a winning strategy to enforce ϕ_2, and the rule is then safe. We come back to that later in the paper (see Lemma 1).

Chatterjee et al. in [6] proposed synthesis rules inspired by Win-under-Hyp but avoid the aforementioned problem. The rule was originally proposed in a model with two components and a scheduler. We study here two natural extensions for n players.

Rules $\mathsf{AG}^{\wedge}$ and $\mathsf{AG}^{\vee}$: The objectives are achieved by

($\mathsf{AG}^{\wedge}$) *assume-guarantee-*$\wedge$ if there exists a strategy profile $\sigma_{\mathcal{P}}$ such that 1. $\mathsf{G}, \sigma_{\mathcal{P}} \models \bigwedge_{i \in \mathcal{P}} \phi_i$, 2. for all players i, $\mathsf{G}, \sigma_i \models (\bigwedge_{j \in \mathcal{P} \setminus \{i\}} \phi_j) \Rightarrow \phi_i$.

($\mathsf{AG}^{\vee}$) *assume-guarantee-*$\vee$[1] if there exists a strategy profile $\sigma_{\mathcal{P}}$ such that 1. $\mathsf{G}, \sigma_{\mathcal{P}} \models \bigwedge_{i \in \mathcal{P}} \phi_i$, 2. for all players i, $\mathsf{G}, \sigma_i \models (\bigvee_{j \in \mathcal{P} \setminus \{i\}} \phi_j) \Rightarrow \phi_i$.

[1] This rule was introduced in [5], under the name Doomsday equilibria, as a generalization of the AG rule of [6] to the case of n-players.

The two rules differ in the second requirement: $\mathsf{AG}^{\wedge}$ requires that player i wins whenever *all* the other players win, while $\mathsf{AG}^{\vee}$ requires player i to win whenever *one* of the other player wins. Clearly $\mathsf{AG}^{\vee}$ is stronger, and the two rules are equivalent for two-player games. As shown in [8], for two-player games, a profile of strategy for $\mathsf{AG}^{\wedge}$ (or $\mathsf{AG}^{\vee}$) is a Nash equilibrium in a derived game where players want, in lexicographic order, first to satisfy their own objectives, and then as a secondary objective, want to falsify the objectives of the other players. As NE, $\mathsf{AG}^{\wedge}$ and $\mathsf{AG}^{\vee}$ require players to *synchronize* on a particular strategy profiles. As we will see, this is not the case for the new rule that we propose.

[16] and [21] introduce two versions of *rational synthesis* (RS). In the two cases, one of the player, say player 1, models the system while the other players model the environment. The existential version ($\mathsf{RS}^{\exists}$) searches for a strategy for the system, and a profile of strategies for the environment, such that the objective of the system is satisfied, and the profile for the environment is *stable* according to a solution concept which is either NE, SPE, or Dom. The universal version ($\mathsf{RS}^{\forall}$) searches for a strategy for the system, such that for all environment strategy profiles that are *stable* according to the solution concept, the objective of the system holds. We write $\Sigma^{\mathsf{NE}}_{G,\sigma_1}$, resp. $\Sigma^{\mathsf{SPE}}_{\mathsf{G},\sigma_1}$, for the set of strategy profiles $\sigma_{-1} = (\sigma_2, \sigma_3, \ldots, \sigma_n)$ that are NE (resp. SPE) equilibria in the game G when player 1 plays σ_1, and $\Sigma^{\mathsf{Dom}}_{G,\sigma_1}$ for the set of strategy profiles σ_{-1} where each strategy σ_j, $2 \leq j \leq n$, is dominant in the game G when player 1 plays σ_1.

Rules $\mathsf{RS}^{\exists,\forall}(\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom})$: Let $\gamma \in \{\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom}\}$, the objective is achieved by:

($\mathsf{RS}^{\exists}(\gamma)$) existential rational synthesis under γ if there is a strategy σ_1 of player 1, and a profile $\sigma_{-1} \in \Sigma^{\gamma}_{\mathsf{G},\sigma_1}$, such that $\mathsf{G}, \sigma_1, \sigma_{-1} \models \phi_1$.

($\mathsf{RS}^{\forall}(\gamma)$) universal rational synthesis under γ if there is a strategy σ_1 of player 1, such that $\Sigma^{\gamma}_{\mathsf{G},\sigma_1} \neq \emptyset$, and for all $\sigma_{-1} \in \Sigma^{\gamma}_{\mathsf{G},\sigma_1}$, $\mathsf{G}, \sigma_1, \sigma_{-1} \models \phi_1$.

Clearly, ($\mathsf{RS}^{\forall}(\gamma)$) is stronger than ($\mathsf{RS}^{\exists}(\gamma)$) and more robust. As $\mathsf{RS}^{\exists,\forall}(\mathsf{NE}, \mathsf{SPE})$ are derived from NE and SPE, they require players to synchronize on particular strategy profiles.

Novel rule We now present our novel rule based on the notion of *admissible strategies*.

Rule AA: The objectives are achieved by *assume-admissible (*AA*) strategies* if there is a strategy profile $\sigma_{\mathcal{P}}$ such that: 1. for all $i \in \mathcal{P}$, $\sigma_i \in \mathsf{Adm}_i(\mathsf{G})$; 2. for all $i \in \mathcal{P}$, $\forall \sigma'_{-i} \in \mathsf{Adm}_{-i}(\mathsf{G})$. $\mathsf{G}, \sigma'_{-i}, \sigma_i \models \phi_i$.

A player-i strategy satisfying conditions 1 and 2 above is called *assume-admissible-winning* (AA-*winning*). A profile of AA-winning strategies is an AA-*winning strategy profile*. The rule AA requires that each player has a strategy *winning* against *admissible* strategies of other players. So we assume that players do not play strategies which are *dominated*, which is reasonable as dominated strategies are clearly *suboptimal options*.

Contrary to Coop, $\mathsf{AG}^{\wedge}$, and $\mathsf{AG}^{\vee}$, AA does not require that the strategy profile is winning for each player. As for Win, this is a consequence of the definition:

Theorem 1. *For all* AA*-winning strategy profile* $\sigma_{\mathcal{P}}$, $\mathsf{G}, \sigma_{\mathcal{P}} \models \bigwedge_{i \in \mathcal{P}} \phi_i$.

The condition that AA strategies are admissible is necessary for Thm. 1; it does not suffice to have strategies that are winning against admissible strategies (see Appendix).

4 Synthesis Rules at the Light of an Example

We illustrate the synthesis rules on an example of a real-time scheduler with two tasks. The system is composed of Sched (player 1) and Env (player 2). Env chooses the truth value for r_1, r_2 (r_i is a request for task i), and Sched controls q_1, q_2 (q_i means that task i has been scheduled). Our model is a turn-based game: first, Env chooses a value for r_1, r_2, then in the next round Sched chooses a value for q_1, q_2, and we repeat forever. The requirements for Sched and Env are as follows: 1. Sched is not allowed to schedule the two tasks at the same time. When r_1 is true, then task 1 must be scheduled (q_1) *within* three rounds. When r_2 is true, task 2 must be scheduled (q_2) in *exactly* three rounds. 2. Whenever Env issues r_i then it does not issue this request again before the occurrence of the grant q_i. Env issues infinitely many requests r_1 and r_2. We say that a request r_i is *pending* whenever the corresponding grant has not yet been issued. Those requirements can be expressed in LTL as follows:

- $\phi_{\mathsf{Sched}} = \mathtt{G}(r_1 \rightarrow \mathtt{X}q_1 \vee \mathtt{XXX}q_1) \wedge \mathtt{G}(r_2 \rightarrow \mathtt{XXX}q_2) \wedge \mathtt{G}\neg(q_1 \wedge q_2)$.
- $\phi_{\mathsf{Env}} = \mathtt{G}(r_1 \rightarrow \mathtt{X}(\neg r_1 \mathtt{W} q_1)) \wedge \mathtt{G}(r_2 \rightarrow \mathtt{X}(\neg r_2 \mathtt{W} q_2)) \wedge (\mathtt{GF}r_1) \wedge (\mathtt{GF}r_2)$.

A solution compatible with the previous rules in the literature. First, we note that there is no winning strategy neither for Sched, nor for Env. In fact, first let $\hat{\sigma_1}$ be the strategy of Sched that never schedules any of the two tasks, i.e. leaves q_1 and q_2 constantly false. This is clearly forcing $\neg\phi_{\mathsf{Env}}$ against all strategies of Env. Second, let $\hat{\sigma_2}$ be s.t. Env always requests the scheduling of both task 1 and task 2, i.e. r_1 and r_2 are constantly true. It is easy to see that this enforces $\neg\phi_{\mathsf{Sched}}$ against any strategy of Sched. So, there is no solution with rule Win^2. But clearly those strategies are also not compatible with the objectives of the respective players, so this leaves the possibility to apply successfully the other rules. We now consider a strategy profile which is a solution for all the rules except for AA.

Let (σ_1, σ_2) be strategies for player 1 and 2 respectively, such that the outcome of (σ_1, σ_2) is "Env emits r_1, then Sched emits q_1, Env emits r_2, then Sched waits one round and emits q_2, and repeat." If a deviation from this *exact* execution is observed, then the two players switch to strategies $\hat{\sigma_1}$ and $\hat{\sigma_2}$ respectively, i.e. to the strategies that falsify the specification of the other players. The reader can now convince himself/herself that (σ_1, σ_2) is a solution for Coop, AG and $\mathsf{RS}^{\exists}(\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom})$. Furthermore, we claim that σ_1 is a solution for Win-under-Hyp and $\mathsf{RS}^{\forall}(\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom})$.[3] But, assume now that Env is a device driver which requests the scheduling of tasks by the scheduler of the kernel of an OS. Clearly (σ_1, σ_2), which is compatible with all the previous rules (but Win), makes little sense in this context. On the other hand, ϕ_{Sched} and ϕ_{Env} are natural specifications for such a system. So, there is clearly room for other synthesis rules!

Solutions provided by AA, our novel rule. For Env, we claim that the set of *admissible strategies*, noted $\mathsf{Adm}(\phi_{\mathsf{Env}})$, are exactly those that (i) do not emit a new request before the previous one has been acknowledged, and (ii) do always eventually emit a (new) request when the previous one has been granted. Indeed as we have seen above, Env and Sched can cooperate to satisfy $\phi_{\mathsf{Sched}} \wedge \phi_{\mathsf{Env}}$, so any strategy of Env which would imply the falsification of ϕ_{Env} is dominated and so it is not admissible. Also, we have seen that Env does not have a winning strategy for ϕ_{Env}, so Env cannot do better.

Now, let us consider the following strategy for Sched. (i) if pending requests r_1 and r_2 were made one round ago, then grant q_1; if pending requests r_1 and r_2 were made three rounds ago, then behave arbitrarily (it is no more possible to satisfy the specification); (ii) if pending request r_2 was made three rounds ago, but not r_1, then grant q_2; (iii) if pending r_1 was made three rounds ago, but not r_2, then grant q_1. We claim that this strategy is *admissible* and while it is *not* winning against *all* possible strategies of Env, it is *winning* against *all admissible* strategies of Env. So, this strategy enforces ϕ_{Sched} against all reasonable strategies of Env w.r.t. to his/her own objective ϕ_{Env}. In fact, there is a whole set of such strategies for Sched, noted $\mathsf{WinAdm}_{\mathsf{Sched}}$. Similarly, there is a whole set of strategies for Env which are both *admissible* and winning against the *admissible* strategies of Sched, noted $\mathsf{WinAdm}_{\mathsf{Env}}$. We prove in the next section that the solutions to AA are *rectangular sets*: they are exactly the solutions in $\mathsf{WinAdm}_{\mathsf{Sched}} \times \mathsf{WinAdm}_{\mathsf{Env}}$. This ensures that AA leads to *resilient* solutions: players do not need to synchronize with the other players on a particular strategy profile but they can arbitrarily choose inside their sets of strategies that are admissible and winning against the admissible strategies of the other players.

5 Comparison of Synthesis Rules

In this section, we compare the synthesis rules to understand which ones yield solutions more often, and to assess their robustness. Some relations are easy to establish; for instance, rules $\mathsf{Win}, \mathsf{AG}^{\vee}, \mathsf{AG}^{\wedge}, \mathsf{AA}$ imply Coop by definition (and Thm. 1). We summarize the implication relations between the rules in Fig. 1. We present the rules $\mathsf{AG}^{\vee}, \mathsf{AG}^{\wedge}$, and the variants of $\mathsf{RS}^{\cdot}(\cdot)$ in one group, respectively. A dashed arrow from A to B means that rule A implies *some* rule in B; while a plain arrow means that A implies *all* rules in B (*e.g.* AA

[2]Also, it is easy to see that Env does not have a dominant strategy for his specification (details are in Appendix). So, considering dominant strategies as *best-effort strategies* would not lead to a solution for this example. To find a solution, we need to take into account the objectives of the other players.

[3]The interested reader can find a detailed explanation for all those claims in the appendix.

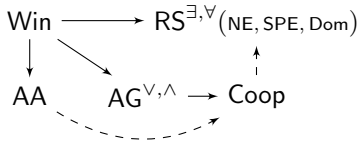

Figure 1: Comparison of synthesis rules.

implies $\mathsf{AG}^{\wedge}$ but not $\mathsf{AG}^{\vee}$; while Win implies both rules.) An absence of path means that A does not imply any variant of B. Thus the figure explains which approaches yield solutions more often, by abstracting away the precise variants. The following theorem states the correctness of our diagram; the appendix contains detailed proofs between all pairs of rules.

Theorem 2. *The implication relations of Fig. 1 hold.*

In the controller synthesis framework using two-player games between a controller and its environment, some works advocate the use of environment objectives which the environment can guarantee against any controller [7]. Under this assumption, Win-under-Hyp implies AA:

Lemma 1. *Let* $\mathsf{G} = \langle \mathsf{A}, \phi_1, \phi_2 \rangle$ *be a two-player game. If player* 2 *has a winning strategy for* ϕ_2 *and* Win-under-Hyp *has a solution, then* AA *has a solution.*

We now consider the *robustness* of the profiles synthesized using the above rules. An AA-winning strategy profile $\sigma_{\mathcal{P}}$ is robust in the following sense: The set of AA-winning profiles is *rectangular*, *i.e. any* combination of AA-winning strategies independently chosen for each player, is an AA-winning profile. Second, if one replaces *any* subset of strategies in AA-winning profile $\sigma_{\mathcal{P}}$ by arbitrary admissible strategies, the objectives of all the other players still hold. Formally, a *rectangular set* of strategy profiles is a set that is a Cartesian product of sets of strategies, given for each player. A synthesis rule is *rectangular* if the set of strategy profiles satisfying the rule is rectangular. The RS rules require a specific definition since player 1 has a particular role: we say that $\mathsf{RS}^{\forall,\exists}(\gamma)$ is rectangular if given for any strategy σ_1 witnessing the rule, the set of strategy profiles $(\sigma_2, \ldots, \sigma_n) \in \Sigma^{\gamma}_{\mathsf{G},\sigma_1}$ such that $\mathsf{G}, \sigma_1, \ldots, \sigma_n \models \phi_1$ is rectangular. We show that apart from AA, only Win and $\mathsf{RS}^{\forall}(\mathsf{Dom})$ are rectangular among the other rules:

Theorem 3. *We have 1. Rule* AA *is rectangular; and for all games* G*,* AA*-winning strategy profile* σ_P*, coalition* $C \subseteq \mathcal{P}$*, if* $\sigma'_C \in \mathsf{Adm}_C(\mathsf{G})$*, then* $\mathsf{G}, \sigma_{-C}, \sigma'_C \models \bigwedge_{i \in -C} \phi_i$*. 2. The rules* Win *and* $\mathsf{RS}^{\forall}(\mathsf{Dom})$ *are rectangular; the rules* Coop*,* $\mathsf{AG}^{\vee}$*,* $\mathsf{AG}^{\wedge}$*,* $\mathsf{RS}^{\exists}(\mathsf{NE},\mathsf{SPE},\mathsf{Dom})$*, and* $\mathsf{RS}^{\forall}(\mathsf{NE},\mathsf{SPE})$ *are not rectangular.*

6 Algorithm for Assume-Admissible Synthesis

In this section, we recall the characterization of the outcomes of admissible strategy profiles given in [4], and we derive algorithms for the AA rule. We use the game of Fig. 2 as a running example for this section. Clearly, none of the players of this game has a winning strategy for his own objective when not taking into account the objective of the other player, but, as we will see, both players have an admissible and winning strategy against the *admissible* strategies of the other player, and so the AA rule applies.

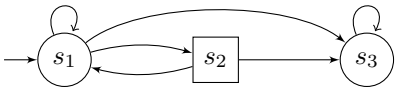

Figure 2: Game G with two players $\mathcal{P} = \{1, 2\}$. Player 1 controls the round states, and has objective $\mathtt{GF}s_2$, and player 2 controls the square state and has objective $\mathtt{GF}s_1$.

The notion of *value* associated to the states of a game plays an important role in the characterization of admissible strategies and their outcomes [2, 4]. Fix a game G. A state s

has value 1 for player i, written $\mathsf{Val}_i(s) = 1$, if player i has a winning strategy from s; $\mathsf{Val}_i(s) = -1$ if for all strategy profiles $\sigma_{\mathcal{P}} \in \Sigma_{\mathcal{P}}$, $\mathsf{Out}_{\mathsf{G},s}(\sigma_{\mathcal{P}})$ does not satisfy ϕ_i; and otherwise $\mathsf{Val}_i(s) = 0$. A *player j decreases its own value in history h* if there is a position k such that $\mathsf{Val}_j(h_k) > \mathsf{Val}_j(h_{k+1})$ and $h_k \in \mathsf{S}_j$. We proved in [4], that admissible strategies do not decrease their own values. Let us call such strategies *value-preserving*. In fact, if the current state has value 1, there is a winning strategy which stays within the winning region; if the value is 0, then although other players may force the play into states of value -1, a good strategy for player i will not do this by itself.

Lemma 2 ([4, Lem. 1])**.** *For all games* G*, players* i*, and histories* ρ*, if* $\mathsf{last}(\rho) \in \mathsf{S}_i$ *and* $\sigma_i \in \mathsf{Adm}_i$ *then* $\mathsf{Val}_i(\delta(\mathsf{last}(\rho), \sigma_i(\rho))) = \mathsf{Val}_i(\mathsf{last}(\rho))$.

For player i, let us define the sets $V_{i,x} = \{s \mid \mathsf{Val}_i(s) = x\}$ for $x \in \{-1, 0, 1\}$, which partition S. We define the set of *value-preserving edges* for player i as $E_i = \{(s, a) \in \mathsf{S} \times \mathsf{Act} \mid s \in \mathsf{S}_i \Rightarrow \mathsf{Val}_i(\delta(s, a)) = \mathsf{Val}_i(s)\}$. Observe that value-preserving strategies for player i are exactly those respecting E_i. In our running example of Fig. 2, it should be clear that any strategy that chooses a transition that goes to s_3 is *not* admissible nor for Player 1 neither for Player 2, as by making this choice both players are condemned to lose their own objective while their other choices leave a chance to win; so the choice of going to s_3 would decrease their own value. So, we can already conclude that Player 2 always chooses $s_2 \mapsto s_1$, his only admissible strategy.

Not all value-preserving strategies are admissible: for Müller objectives, staying inside the winning region does not imply the objective. Moreover, in states of value 0, admissible strategies must visit states where other players can "help" satisfy the objective. Formally, *help states* for player i are other players' states with value 0 and at least two different successors of value 0 or 1. Let $H_i = \{s \in \mathsf{S} \setminus \mathsf{S}_i \mid \mathsf{Val}_i(s) = 0 \wedge \exists s' \neq s''.\ s' \in \delta(s, \mathsf{Act}) \wedge s'' \in \delta(s, \mathsf{Act}) \wedge \mathsf{Val}_i(s') \geq 0 \wedge \mathsf{Val}_i(s'') \geq 0\}$. Given this, the following lemma, adapted from [4], characterizes the outcomes of admissible strategies. We denote by $\mathtt{G}(E_i)$ the set of runs that respect E_i, *i.e.* $\mathtt{G}(\bigvee_{(s,a)\in E_i} s \wedge \mathtt{X}(\delta(s, a)))$.

Lemma 3. *For all games* G*, and players* i*,* $\mathsf{Out}_{\mathsf{G}} \cap \Phi_i = \mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_i, \Sigma_{-i})$*, where* $\Phi_i = \mathtt{G}(E_i) \wedge (\mathtt{GF}(V_{i,1}) \Rightarrow \phi_i) \wedge (\mathtt{GF}(V_{i,0}) \Rightarrow \phi_i \vee \mathtt{GF}(H_i))$.

In our running example of Fig. 2, a strategy of Player 1 which, after some point, always chooses $s_1 \mapsto s_1$ is dominated by strategies that chose infinitely often $s_1 \mapsto s_2$. This is a corollary of the lemma above. Indeed, while all those strategies only visit states with value 0 (and so do not decrease the value for Player 1), the strategy that always chooses $s_1 \mapsto s_1$ has an outcome which is loosing for Player 1 while the other strategies are compatible with outcomes that are winning for Player 1. So, outcome of admissible strategies for Player 1 that always visit states with values 0, also visits s_2 infinitely often. Using the fact that strategies are value-preserving and the last observation, we can now conclude that both players have (admissible) winning strategies against the admissible strategies of the other players. For instance when Player 1 always chooses to play $s_1 \mapsto s_2$, he wins against the admissible strategies of Player 2.

Note that Φ_i can be decomposed into a safety condition $S_i = \mathtt{G}(E_i)$ and a *prefix independent* condition $M_i = (\mathtt{GF}(V_{i,1}) \Rightarrow \phi_i) \wedge (\mathtt{GF}(V_{i,0}) \Rightarrow (\phi_i \vee \mathtt{GF}(H_i))$ which can be expressed by a Müller condition described by a circuit of polynomial size.

For player i, we let $\Omega_i = \mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_i) \wedge (\mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_{-i}) \Rightarrow \phi_i)$, which describes the outcomes of admissible strategies of player i, which satisfy objective ϕ_i under the hypothesis that they are compatible with other players' admissible strategies. In fact, it follows from [4] that Ω_i captures the *outcomes* of AA-winning strategies for player i.

Lemma 4. *A player i strategy is* AA*-winning iff it is winning for objective* Ω_i*.*

Objective Ω_i is not directly expressible as a Müller condition, since Φ_i and $\bigwedge_j \Phi_j$ contain safety parts. Nevertheless, the information whether $\mathtt{G}(E_i)$, or $\mathtt{G}(\cup_{j\neq i} E_j)$ has been violated can be encoded in the state space. Formally, for each player i, we define game G'_i by taking the product of G with $\{\top, 0, \bot\}$; that is, the states are $\mathsf{S} \times \{\top, 0, \bot\}$, and the initial state $(s_{\mathsf{init}}, 0)$. The transitions are defined as for G for the first component; while from state $(s, 0)$, any action a outside E_i leads to $(\delta(s, a), \bot)$, and any action a outside E_j, $j \neq i$, leads to $(\delta(s, a), \top)$. The second component is absorbing at $\bot, \top$. We now rewrite the condition Ω_i

for G'_i as $\Omega'_i = \big(\mathtt{GF}(\mathsf{S} \times \{0\}) \wedge M'_i \wedge (\wedge_{j \neq i} M'_j \Rightarrow \phi'_i)\big) \vee (\mathtt{GF}(\mathsf{S} \times \{\top\}) \wedge M'_i)$, where M'_i is the set of runs of G'_i whose projections to G are in M_i, and similarly for ϕ'_i.

Now, checking AA-synthesis is reduced to solving games with Müller conditions. Moreover, we also obtain a polynomial-time algorithm when all objectives are Büchi conditions, by showing that Ω'_i is expressible by a parity condition with four colors.

Theorem 4. AA*-synthesis in multiplayer games is* PSPACE*-complete, and* P*-complete for Büchi objectives. Player i wins for objective Ω_i in G iff he wins for objective Ω'_i in G'_i.*

7 Abstraction

We present abstraction techniques to compute assume-admissible strategy profiles following the *abstract interpretation* framework [11]; see [18] for games. Abstraction is a crucial feature for scalability in practice, and we show here that the AA rule is amenable to abstraction techniques. The problem is not directly reducible to computing AA-winning strategies in abstract games obtained as *e.g.* in [13]; in fact, it can be easily seen that the set of admissible strategies of an abstract game is incomparable with those of the concrete game in general.

Overview. Informally, to compute an AA-winning strategy for player i, we construct an abstract game $\mathcal{A}'_i$ with objective $\underline{\Omega'_i}$ s.t. winning strategies of player i in $\mathcal{A}'_i$ map to AA-winning strategies in G. To define $\mathcal{A}'$, we re-visit the steps of the algorithm of Section 6 by defining approximations computed on the abstract state space. More precisely, we show how to compute under- and over-approximations of the sets $V_{x,k}$, namely $\underline{V}_{x,k}$ and $\overline{V}_{x,k}$, using fixpoint computations on the abstract state space only. We then use these sets to define approximations of the value preserving edges ($\underline{E}_k$ and $\overline{E}_k$) and those of the help states ($\underline{H}_k$ and $\overline{H}_k$). These are then combined to define objective $\underline{\Omega'}_k$ s.t. if player k wins the abstract game for $\underline{\Omega'}_k$, then he wins the original game for Ω'_k, and thus has an AA-winning strategy.

Abstract Games. Consider $\mathsf{G} = \langle \mathsf{A}, (\phi_i)_{i \in \mathcal{P}} \rangle$ with $\mathsf{A} = \langle \mathcal{P}, (\mathsf{S}_i)_{i \in \mathcal{P}}, s_{\mathsf{init}}, (\mathsf{Act}_i)_{i \in \mathcal{P}}, \delta \rangle$ where each ϕ_i is a Müller objective given by a family of sets of states $(\mathcal{F}_i)_{i \in \mathcal{P}}$. Let $\mathsf{S}^{\mathsf{a}} = \biguplus_{i \in \mathcal{P}} \mathsf{S}^{\mathsf{a}}_i$ denote a finite set, namely the *abstract state space*. A *concretization function* $\gamma \colon \mathsf{S}^{\mathsf{a}} \mapsto 2^{\mathsf{S}}$ is a function such that: 1. the abstract states partitions the state space: $\biguplus_{s^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}}} \gamma(s^{\mathsf{a}}) = \mathsf{S}$, 2. it is compatible with players' states: for all players i and $s^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}}_i$, $\gamma(s^{\mathsf{a}}) \subseteq \mathsf{S}_i$. We define the corresponding *abstraction* function $\alpha : \mathsf{S} \to \mathsf{S}^{\mathsf{a}}$ where $\alpha(s)$ is the unique state s^{a} s.t. $s \in \gamma(s^{\mathsf{a}})$. We also extend α, γ naturally to sets of states; and to histories, by replacing each element of the sequence by its image.

We further assume that γ is *compatible* with all objectives $\mathcal{F}_i$ in the sense that the abstraction of a set S is sufficient to determine whether $S \in \mathcal{F}_i$: for all $i \in \mathcal{P}$, for all $S, S' \subseteq \mathsf{S}$ with $\alpha(S) = \alpha(S')$, we have $S \in \mathcal{F}_i \Leftrightarrow S' \in \mathcal{F}_i$. If the objective ϕ_i is given by a circuit, then the circuit for the corresponding abstract objective ϕ^{a}_i is obtained by replacing each input on state s by $\alpha(s)$. We thus have $\rho \in \phi_i$ if, and only if, $\alpha(\rho) \in \phi^{\mathsf{a}}_i$.

The *abstract transition relation* Δ^{a} induced by γ is defined by: $(s^{\mathsf{a}}, a, t^{\mathsf{a}}) \in \Delta^{\mathsf{a}} \Leftrightarrow \exists s \in \gamma(s^{\mathsf{a}}), \exists t \in \gamma(t^{\mathsf{a}}), t = \delta(s, a)$. We write $\mathsf{post}_{\Delta}(s^{\mathsf{a}}, a) = \{t^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}} \mid \Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})\}$, and $\mathsf{post}_{\Delta}(s^{\mathsf{a}}, \mathsf{Act}) = \cup_{a \in \mathsf{Act}} \mathsf{post}_{\Delta}(s^{\mathsf{a}}, a)$. For each coalition $C \subseteq \mathcal{P}$, we define a game in which players C play together against coalition $-C$; and the former resolves non-determinism in Δ^{a}. Intuitively, the winning region for C in this abstract game will be an over-approximation of the original winning region. Given C, the *abstract arena* $\mathcal{A}^C$ is $\langle \{C, -C\}, (\mathsf{S}_C, \mathsf{S}_{-C}), \alpha(s_{\mathsf{init}}),$
$(\mathsf{Act}_C, \mathsf{Act}_{-C}), \delta^{\mathsf{a},C} \rangle$, where $\mathsf{S}_C = \big(\bigcup_{i \in C} \mathsf{S}^{\mathsf{a}}_i\big) \cup \big(\bigcup_{i \in \mathcal{P}} \mathsf{S}^{\mathsf{a}}_i \times \mathsf{Act}_i\big)$, $\mathsf{S}_{-C} = \bigcup_{i \notin C} \mathsf{S}^{\mathsf{a}}_i$; and $\mathsf{Act}_C = \big(\bigcup_{i \in C} \mathsf{Act}_i\big) \cup \mathsf{S}^{\mathsf{a}}$ and $\mathsf{Act}_{-C} = \bigcup_{i \in -C} \mathsf{Act}_i$. The relation $\delta^{\mathsf{a},C}$ is given by: if $s^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}}$, then $\delta^{\mathsf{a},C}(s^{\mathsf{a}}, a) = (s^{\mathsf{a}}, a)$. If $(s^{\mathsf{a}}, a) \in \mathsf{S}^{\mathsf{a}} \times \mathsf{Act}$ and $t^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}}$ satisfies $(s^{\mathsf{a}}, a, t^{\mathsf{a}}) \in \Delta^{\mathsf{a}}$ then $\delta^{\mathsf{a},C}((s^{\mathsf{a}}, a), t^{\mathsf{a}}) = t^{\mathsf{a}}$; while for $(s^{\mathsf{a}}, a, t^{\mathsf{a}}) \notin \Delta^{\mathsf{a}}$, the play leads to an arbitrarily chosen state u^{a} with $\Delta(s^{\mathsf{a}}, a, u^{\mathsf{a}})$. Thus, from states (s^{a}, a), coalition C chooses a successor t^{a}.

We extend γ to histories of $\mathcal{A}^C$ by first removing states of $(\mathsf{S}^{\mathsf{a}}_i \times \mathsf{Act}_i)$; and extend α by inserting these intermediate states. Given a strategy σ of player k in $\mathcal{A}^C$, we define its *concretization* as the strategy $\gamma(\sigma)$

of G that, at any history h of G, plays $\gamma(\sigma)(h) = \sigma(\alpha(h))$. We write $\mathsf{Win}_D(\mathcal{A}^C, \phi^{\mathsf{a}}_k)$ for the states of S^{a} from which the coalition D has a winning strategy in $\mathcal{A}^C$ for objective ϕ^{a}_k, with $D \in \{C, -C\}$. Informally, it is easier for coalition C to achieve an objective in $\mathcal{A}^C$ than in G, that is, $\mathsf{Win}_C(\mathcal{A}^C, \phi^{\mathsf{a}}_k)$ over-approximates $\mathsf{Win}_C(\mathsf{A}, \phi_k)$:

Lemma 5. *If the coalition C has a winning strategy for objective ϕ_k in G from s then it has a winning strategy for ϕ^{a}_k in $\mathcal{A}^C$ from $\alpha(s)$.*

Value-Preserving Strategies. We now provide under- and over-approximations for value-preserving strategies for a given player. We start by computing approximations $\underline{V}_{k,x}$ and $\overline{V}_{k,x}$ of the sets $V_{k,x}$, and then use these to obtain approximations of the value-preserving edges E_k. Fix a game G, and a player k. Let us define the *controllable predecessors* for player k as $\mathsf{CPRE}_{\mathcal{A}^{\mathcal{P}\setminus\{k\}},k}(X) = \{s^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}}_k \mid \exists a \in \mathsf{Act}_k, \mathsf{post}_\Delta(s^{\mathsf{a}}, a) \subseteq X\} \cup \{s^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}}_{\mathcal{P}\setminus\{k\}} \mid \forall a \in \mathsf{Act}_{-k}, \mathsf{post}_\Delta(s^{\mathsf{a}}, a) \subseteq X\}$. We let

$$
\begin{array}{ll}
\overline{V}_{k,1} = \mathsf{Win}_{\{k\}}(\mathcal{A}^{\{k\}}, \phi^{\mathsf{a}}_k), & \overline{V}_{k,-1} = \mathsf{Win}_{\emptyset}(\mathcal{A}^{\emptyset}, \neg\phi^{\mathsf{a}}_k),\\
\overline{V}_{k,0} = \mathsf{Win}_{\mathcal{P}\setminus\{k\}}(\mathcal{A}^{\mathcal{P}\setminus\{k\}}, \neg\phi^{\mathsf{a},k}) \cap \mathsf{Win}_{\mathcal{P}}(\mathcal{A}^{\mathcal{P}}, \phi^{\mathsf{a}}_k), & \\
\underline{V}_{k,1} = \mathsf{Win}_{\{k\}}(\mathcal{A}^{\mathcal{P}\setminus\{k\}}, \phi^{\mathsf{a}}_k), & \underline{V}_{k,-1} = \mathsf{Win}_{\emptyset}(\mathcal{A}^{\mathcal{P}}, \neg\phi^{\mathsf{a}}_k)\\
\underline{V}_{k,0} = \nu X.\big(\mathsf{CPRE}_{\mathcal{A}^{\mathcal{P}\setminus\{k\}},k}(X \cup \underline{V}_{k,1} \cup \underline{V}_{k,-1}) \cap F\big), & \\
\quad \text{where } F = \mathsf{Win}_{\mathcal{P}\setminus\{k\}}(\mathcal{A}^{\{k\}}, \neg\phi^{\mathsf{a}}_k) \cap \mathsf{Win}_{\mathcal{P}}(\mathcal{A}^{\emptyset}, \phi^{\mathsf{a}}_k). &
\end{array}
$$

The last definition uses the $\nu X.f(X)$ operator which is the greatest fixpoint of f. These sets define approximations of the sets $V_{k,x}$. Informally, this follows from the fact that to define *e.g.* $\overline{V}_{k,1}$, we use the game $\mathcal{A}^{\{k\}}$, where player k resolves itself the non-determinism, and thus has more power than in G. In contrast, for $\underline{V}_{k,1}$, we solve $\mathcal{A}^{\mathcal{P}\setminus\{k\}}$ where the adversary resolves non-determinism. We state these properties formally:

Lemma 6. *For all players k and $x \in \{-1, 0, 1\}$, $\gamma(\underline{V}_{k,x}) \subseteq V_{k,x} \subseteq \gamma(\overline{V}_{k,x})$.*

We thus have $\cup_x \gamma(\overline{V}_{k,x}) = \mathsf{S}$ (as $\cup_x V_{k,x} = \mathsf{S}$) but this is not the case for $\underline{V}_{k,x}$; so let us define $\underline{V} = \cup_{j\in\{-1,0,1\}} \underline{V}_{k,j}$. We now define approximations of E_k based on the above sets.

$$
\begin{array}{l}
\overline{E}_k = \{(s^{\mathsf{a}}, a) \in \mathsf{S}^{\mathsf{a}} \times \mathsf{Act} \mid s^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}}_k \Rightarrow \exists x, s^{\mathsf{a}} \in \overline{V}_{k,x}, \mathsf{post}_\Delta(s^{\mathsf{a}}, a) \cap \cup_{l\geq x} \overline{V}_{k,l} \neq \emptyset\},\\
\underline{E}_k = \{(s^{\mathsf{a}}, a) \in \mathsf{S}^{\mathsf{a}} \times \mathsf{Act} \mid s^{\mathsf{a}} \in \mathsf{S}^{\mathsf{a}}_k \Rightarrow \exists x, s^{\mathsf{a}} \in \underline{V}_{k,x}, \mathsf{post}_\Delta(s^{\mathsf{a}}, a) \subseteq \cup_{l\geq x} \underline{V}_{k,l}\} \cup \{(s^{\mathsf{a}}, a) \mid s^{\mathsf{a}} \notin \underline{V}\}.
\end{array}
$$

Intuitively, $\overline{E}_k$ is an over-approximation of E_k, and $\underline{E}_k$ under-approximates E_k when restricted to states in $\underline{V}$ (notice that $\underline{E}_k$ contains all actions from states outside $\underline{V}$). In fact, our under-approximation will be valid only inside $\underline{V}$; but we will require the initial state to be in this set, and make sure the play stays within $\underline{V}$. We show that sets $\underline{E}_k$ and $\overline{E}_k$ provide approximations of value-preserving strategies.

Lemma 7. *For all games G, and players k, $\mathsf{Strat}_k(E_k) \subseteq \gamma(\mathsf{Strat}_k(\overline{E}_k))$, and if $s_{\mathsf{init}} \in \gamma(\underline{V})$, then $\emptyset \neq \gamma(\mathsf{Strat}_k(\underline{E}_k)) \subseteq \mathsf{Strat}_k(E_k)$.*

Abstract Synthesis of AA-winning strategies. We now describe the computation of AA-winning strategies in abstract games. Consider game G and assume sets $\underline{E}_i, \overline{E}_i$ are computed for all players i. Roughly, to compute a strategy for player k, we will constrain him to play only edges from $\underline{E}_k$, while other players j will play in $\overline{E}_j$. By Lemma 7, any strategy of player k maps to value-preserving strategies in the original game, and all value-preserving strategies for other players are still present. We now formalize this idea, incorporating the help states in the abstraction.

We fix a player k. We construct an abstract game in which winning for player k implies that player k has an effective AA-winning strategy in G. We also define $\mathcal{A}'_k = \langle\{\{k\}, -k\}, (\mathsf{S}'^{\mathsf{a}}_k, \mathsf{S}'^{\mathsf{a}}_{-k} \cup \mathsf{S}'^{\mathsf{a}} \times \mathsf{Act}), \alpha(s_{\mathsf{init}}), (\mathsf{Act}_k, \mathsf{Act}_{-k}), \delta_{\mathcal{A}^k}\rangle$, where $\mathsf{S}'^{\mathsf{a}} = \mathsf{S}^{\mathsf{a}} \times \{\bot, 0, \top\}$; thus we modify $\mathcal{A}^{\mathcal{P}\setminus\{k\}}$ by taking the product of the state space with $\{\top, 0, \bot\}$. Intuitively, as in Section 6, initially the second component is 0, meaning that no player has violated the value-preserving edges. The component becomes $\bot$ whenever player k plays an action outside of $\underline{E}_k$; and $\top$ if another player j plays outside $\overline{E}_j$. We extend γ to $\mathcal{A}'_k$ by $\gamma((s^{\mathsf{a}}, x)) = \gamma(s^{\mathsf{a}}) \times \{x\}$, and extend it to

histories of $\mathcal{A}'_k$ by first removing the intermediate states $\mathsf{S'}^{\mathsf{a}} \times \mathsf{Act}$. We thus see $\mathcal{A}'_k$ as an abstraction of A' of Section 6.

In order to define the objective of $\mathcal{A}'_k$, let us first define approximations of the help states H_k, where we write $\Delta(s^{\mathsf{a}}, \mathsf{Act}, t^{\mathsf{a}})$ to mean $\exists a \in \mathsf{Act}, \Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})$.

$$\begin{aligned}\overline{H}_k &= \{s^{\mathsf{a}} \in \overline{V}_{k,0} \setminus \mathsf{S}^{\mathsf{a}}_k \mid \exists t^{\mathsf{a}}, u^{\mathsf{a}} \in \overline{V}_{k,0} \cup \overline{V}_{k,1}.\ \Delta(s^{\mathsf{a}}, \mathsf{Act}, t^{\mathsf{a}}) \wedge \Delta(s^{\mathsf{a}}, \mathsf{Act}, u^{\mathsf{a}})\} \\ \underline{H}_k &= \{s^{\mathsf{a}} \in \underline{V}_{k,0} \setminus \mathsf{S}^{\mathsf{a}}_k \mid \exists a \neq b \in \mathsf{Act}, \mathsf{post}_\Delta(s^{\mathsf{a}}, a) \cap \mathsf{post}_\Delta(s^{\mathsf{a}}, b) = \emptyset, \\ &\qquad\qquad \mathsf{post}_\Delta(s^{\mathsf{a}}, a) \cup \mathsf{post}_\Delta(s^{\mathsf{a}}, b) \subseteq \underline{V}_{k,0} \cup \underline{V}_{k,1}\}.\end{aligned}$$

We define the following approximations of the objectives M'_k and Ω'_k in $\mathcal{A}'_k$.

$$\begin{aligned}\underline{M'_k} &= (\mathtt{GF}(\overline{V}_{k,1}) \Rightarrow \phi^{\mathsf{a}}_k) \wedge \left(\mathtt{GF}(\overline{V}_{k,0}) \Rightarrow (\phi^{\mathsf{a}}_k \vee \mathtt{GF}(\underline{H}_k))\right), \\ \overline{M'_k} &= (\mathtt{GF}(\underline{V}_{k,1}) \Rightarrow \phi^{\mathsf{a}}_k) \wedge \left(\mathtt{GF}(\underline{V}_{k,0}) \Rightarrow (\phi^{\mathsf{a}}_k \vee \mathtt{GF}(\overline{H}_k))\right), \\ \underline{\Omega'_k} &= \left(\mathtt{GF}(\mathsf{S}^{\mathsf{a}} \times \{0\}) \wedge \underline{M'_k} \wedge \left(\textstyle\bigwedge_{j \neq k} \overline{M'}_j \Rightarrow \phi^{\mathsf{a}}_k\right)\right) \vee (\mathtt{GF}(\mathsf{S}^{\mathsf{a}} \times \{\top\}) \wedge \underline{M'_k}).\end{aligned}$$

Theorem 5. *For all games G, and players k, if $s_{\mathsf{init}} \in \underline{V}$, and player k has a winning strategy in $\mathcal{A}'_k$ for objective $\underline{\Omega}'_k$, then he has a winning strategy in G'_k for Ω_k; and thus a AA-winning strategy in G.*

Now, if Theorem 5 succeeds to find an AA-winning strategy for each player k, then the resulting strategy profile is AA-winning.

8 Conclusion

In this paper, we have introduced a novel synthesis rule, called the *assume admissible synthesis*, for the synthesis of strategies in non-zero sum n players games played on graphs with omega-regular objectives. We use the notion of admissible strategy, a classical concept from game theory, to take into account the objectives of the other players when looking for winning strategy of one player. We have compared our approach with other approaches such as assume guarantee synthesis and rational synthesis that target the similar scientific objectives. We have developed worst-case optimal algorithms to handle our synthesis rule as well as dedicated abstraction techniques. As future works, we plan to develop a tool prototype to support our assume admissible synthesis rule.

References

[1] Brandenburger Adam, Friedenberg Amanda, H Jerome, et al. Admissibility in games. *Econometrica*, 2008.

[2] Dietmar Berwanger. Admissibility in infinite games. In *Proc. of STACS'07*, volume 4393 of *LNCS*, pages 188–199. Springer, February 2007.

[3] Roderick Bloem, Rüdiger Ehlers, Swen Jacobs, and Robert Könighofer. How to handle assumptions in synthesis. In *SYNT'14*, volume 157 of *EPTCS*, pages 34–50, 2014.

[4] Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS '14, 2014*. ACM, 2014.

[5] Krishnendu Chatterjee, Laurent Doyen, Emmanuel Filiot, and Jean-François Raskin. Doomsday equilibria for omega-regular games. In *VMCAI'14*, volume 8318, pages 78–97. Springer, 2014.

[6] Krishnendu Chatterjee and Thomas A Henzinger. Assume-guarantee synthesis. In *TACAS'07*, volume 4424 of *LNCS*. Springer, 2007.

[7] Krishnendu Chatterjee, Thomas A. Henzinger, and Barbara Jobstmann. Environment assumptions for synthesis. In *CONCUR 2008*, volume 5201 of *LNCS*, pages 147–161. Springer, 2008.

[8] Krishnendu Chatterjee, Thomas A Henzinger, and Marcin Jurdziński. Games with secure equilibria. *Theoretical Computer Science*, 365(1):67–82, 2006.

[9] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.

[10] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronization skeletons using branching-time temporal logic. In *Logics of Programs*, volume 131 of *LNCS*, pages 52–71. Springer, 1981.

[11] Patrick Cousot and Radhia Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL'77*. ACM, 1977.

[12] Werner Damm and Bernd Finkbeiner. Automatic compositional synthesis of distributed systems. In *FM 2014*, volume 8442 of *LNCS*, pages 179–193. Springer, 2014.

[13] Luca de Alfaro, Patrice Godefroid, and Radha Jagadeesan. Three-valued abstractions of games: Uncertainty, but with precision. In *LICS'04*. IEEE, 2004.

[14] E Allen Emerson. Temporal and modal logic. *Handbook of Theoretical Computer Science, Volume B: Formal Models and Sematics (B)*, 995:1072, 1990.

[15] Marco Faella. Admissible strategies in infinite games over graphs. In *MFCS 2009*, volume 5734 of *Lecture Notes in Computer Science*, pages 307–318. Springer, 2009.

[16] Dana Fisman, Orna Kupferman, and Yoad Lustig. Rational synthesis. In *TACAS'10*, volume 6015 of *LNCS*, pages 190–204. Springer, 2010.

[17] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991. Translated into Chinesse by Renin University Press, Bejing: China.

[18] Thomas A. Henzinger, Rupak Majumdar, Freddy Y. C. Mang, and Jean-François Raskin. Abstract interpretation of game properties. In *SAS*, pages 220–239, 2000.

[19] Paul Hunter. *Complexity and Infinite Games on Finite Graphs*. PhD thesis, Computer Laboratory, University of Cambridge, 2007.

[20] Paul Hunter and Anuj Dawar. Complexity bounds for regular games. In Joanna Jedrzejowicz and Andrzej Szepietowski, editors, *MFCS'05*, volume 3618 of *LNCS*. Springer, 2005.

[21] O. Kupferman, G. Perelli, and M.Y. Vardi. Synthesis with rational environments. In *Proc. 12th European Conference on Multi-Agent Systems*, LNCS. Springer, 2014.

[22] David E Long, Anca Browne, Edmund M Clarke, Somesh Jha, and Wilfredo R Marrero. An improved algorithm for the evaluation of fixpoint expressions. In *CAV'94*, volume 818 of *LNCS*. Springer, 1994.

[23] Zohar Manna and Pierre Wolper. Synthesis of communicating processes from temporal logic specifications. In *Logics of Programs*, volume 131 of *LNCS*, pages 253–281. Springer, 1981.

[24] Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi. Reasoning about strategies. In *FSTTCS 2010*, volume 8 of *LIPIcs*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010.

[25] John Nash. Equilibrium points in n-person games. *Proc. NAS*, 1950.

[26] Daniel Neider, Roman Rabinovich, and Martin Zimmermann. Down the borel hierarchy: Solving muller games via safety games. *Theoretical Computer Science*, 560:219–234, 2014.

[27] Amir Pnueli. The temporal logic of programs. In *Foundations of Computer Science, 1977., 18th Annual Symposium on*, pages 46–57. IEEE, 1977.

[28] Helmut Seidl. Fast and simple nested fixpoints. *Inf. Proc. Letters*, 59(6):303–308, 1996.

[29] Michael Ummels. *Stochastic multiplayer games: theory and algorithms*. Amsterdam University Press, 2010.

A Complements on Synthesis Rules (Section 3)

We give the proof of Theorem 1.

Theorem 1. *For all* AA*-winning strategy profile* $\sigma_{\mathcal{P}}$, $\mathsf{G}, \sigma_{\mathcal{P}} \models \bigwedge_{i \in \mathcal{P}} \phi_i$.

Proof. Let $\sigma_{\mathcal{P}}$ be a strategy profile witness of AA. Let i be a player, we have that $\sigma_{-i} \in \mathsf{Adm}_{-i}(\mathsf{G})$, because by condition 1, for all $j \neq i$, $\sigma_j \in \mathsf{Adm}_j(\mathsf{G})$. Then by condition 2 we have that $\mathsf{G}, \sigma_{\mathcal{P}} \models \phi_i$. Since this is true for all players i, we have that $\mathsf{G}, \sigma_{\mathcal{P}} \models \bigwedge_{i \in \mathcal{P}} \phi_i$. □

The following example shows that AA-winning strategies must be admissible themselves for Thm. 1 to hold.

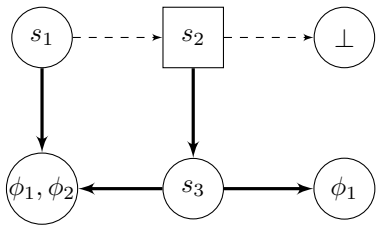

Figure 3: Illustration of the necessity of Condition 1 in the definition of Assume-Admissible Synthesis. Player 1 controls circles and player 2 squares.

Example 1. *In* AA*, the profile of strategy must be composed of admissible strategies only. This is necessary otherwise the players may make assumptions on each other which are not satisfied. This is illustrated by the example of Figure 3 in which two players have reachability objectives* ϕ_1 *and* ϕ_2 *respectively. The states labeled by* ϕ_i *are the target states for the corresponding objective. Admissible strategies are shown in thick edges. Now, the player* 2 *strategy that chooses the dashed edge from* s_2 *satisfies Condition 2 of* AA*, since* s_2 *is not reachable under admissible strategies of player* 1*. Similarly, the player* 1 *strategy that chooses the dashed edge from* s_1 *satisfies Condition 2 of* AA *since the thick edges lead back to a state satisfying* ϕ_1*. But then the resulting profile is such that none of the two players win.*

We explain why the pair of strategies (σ_1, σ_2) is a solution for the synthesis rules Coop, AG, $\mathsf{RS}^{\exists}(\mathsf{NE}, \mathsf{SPE})$:

- Rule Coop: the outcome of (σ_1, σ_2) is "Env emits r_1, then Sched emits q_1, Env emits r_2, then Sched waits one round and emits q_2, and repeat." This outcome is a model for both ϕ_{Sched} and ϕ_{Env}.
- Rules $\mathsf{AG}^{\wedge,\vee}$: the two rules coincide as we have only two players. When both players follow (σ_1, σ_2), we know that the outcome is a model for both ϕ_{Sched} and ϕ_{Env}. We must in addition verify that $\mathsf{Out}(\sigma_1) \models \phi_{\mathsf{Env}} \rightarrow \phi_{\mathsf{Sched}}$ and that $\mathsf{Out}(\sigma_2) \models \phi_{\mathsf{Sched}} \rightarrow \phi_{\mathsf{Env}}$. This is the case as either the outcome of the game between the two players is "Env emits r_1, then Sched emits q_1, Env emits r_2, then Sched waits one round and emits q_2, and repeat", or it is not but then player 1 switches to strategy $\hat{\sigma}_1$ and this forces outcomes that falsify ϕ_{Env}. The argument is similar to justify that $\mathsf{Out}(\sigma_2) \models \phi_{\mathsf{Sched}} \rightarrow \phi_{\mathsf{Env}}$.
- Rules $\mathsf{RS}^{\exists}(\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom})$: for the existential version of rational synthesis, we consider that Sched (player 1) plays the role of the system. We must then prove that (σ_1, σ_2) is such that ϕ_{Sched} is satisfied and σ_2 is stable for NE,SPE, and Dom. This is the case because $\mathsf{Out}(\sigma_1, \sigma_2)$ is a model for ϕ_{Sched}, so it satisfies the specification of the scheduler. Furthermore, it satisfies the specification of the environment, and so the environment has no incentive to deviate: σ_2 is stable for NE. It is also stable for SPE as any deviation of the environment triggers strategy $\hat{\sigma}_1$ for the scheduler and $\hat{\sigma}_2$ for the environment, as a consequence, there is no incentive for environment to deviate from $\hat{\sigma}_2$ as $\mathsf{Out}(\hat{\sigma}_1) \models \neg\phi_{\mathsf{Env}}$. It is stable for Dom because σ_2 is winning under σ_1.

We explain why the pair of strategies σ_1 is a solution for the rules Win-under-Hyp and $\mathsf{RS}^{\forall}(\mathsf{NE}, \mathsf{SPE})$:

- Rule Win-under-Hyp: we have seen above that $\mathsf{Out}(\sigma_1) \models \phi_{\mathsf{Env}} \rightarrow \phi_{\mathsf{Sched}}$ and so σ_1 is a solution for Win-under-Hyp.
- Rules $\mathsf{RS}^{\forall}(\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom})$: against σ_1, the only strategies for the environment that are stable for NE, SPE, Dom are those that make sure that the outcome is equal to "Env emits r_1, then Sched emits q_1, Env emits r_2, then Sched waits one round and emits q_2, and repeat." Indeed, any strategy that would deviate from that outcome would trigger the execution of $\hat{\sigma}_1$ for the scheduler and the outcome would then be losing for the environment making a deviation toward σ_2 profitable for instance.

Absence of dominant strategies Let us now convince the reader that Env has no dominant strategy for his specification. Indeed, it is easy to see that if there are dominant strategies then all admissible strategies must be dominant. Now, take two strategies in $\mathsf{Adm}(\phi_{\mathsf{Env}})$ that behaves the same with the exception that the first strategy when a request is granted (q_i is true), directly issue another request r_i, for $i \in \{1, 2\}$, while the second strategy waits one additional turn to issue the new request r_i. It is easy to see that none of the two strategies dominate the other one, as the Sched can adapt his own strategy to make the first strategy win and the second loose, and the other way around.

B Comparison of the Synthesis Rules (Section 5)

In this section, we detail the comparison of the synthesis notions we consider. We adopt the following notation: an implication $\mathsf{A} \Rightarrow \mathsf{B}$ between two notions A and B means that B has a solution whenever A has one. Note that this does not always imply inclusion between the witnessing strategy profiles. The proof of Theorem 2, will be decomposed in several lemmas.

B.1 Proof of Theorem 2

Remark 1. *We have* $\mathsf{RS}^{\exists}(\mathsf{SPE}) \Rightarrow \mathsf{RS}^{\exists}(\mathsf{NE})$ *and* $\mathsf{RS}^{\forall}(\mathsf{NE}) \Rightarrow \mathsf{RS}^{\forall}(\mathsf{SPE})$ *because any subgame perfect equilibrium is also a Nash equilibrium. Moreover, in the definition of the rules* RS*, the conditions for* $\mathsf{RS}^{\forall}$ *are stronger than for* $\mathsf{RS}^{\exists}$*, so* $\mathsf{RS}^{\forall}(\mathsf{SPE}) \Rightarrow \mathsf{RS}^{\exists}(\mathsf{SPE})$*,* $\mathsf{RS}^{\forall}(\mathsf{NE}) \Rightarrow \mathsf{RS}^{\exists}(\mathsf{NE})$ *and* $\mathsf{RS}^{\forall}(\mathsf{Dom}) \Rightarrow \mathsf{RS}^{\exists}(\mathsf{Dom})$*.*

Lemma 8. $\mathsf{Win} \Rightarrow \mathsf{AA} \Rightarrow \mathsf{Coop} \Rightarrow \mathsf{RS}^{\exists}(\mathsf{SPE})$ *and* $\mathsf{Win} \Rightarrow \mathsf{AG}^{\vee} \Rightarrow \mathsf{AG}^{\wedge} \Rightarrow \mathsf{Coop}$*,*

Proof. $\boxed{\mathsf{Win} \Rightarrow \mathsf{AA}}$ This holds because winning strategies are always admissible [2], therefore a profile witness of Win satisfies condition 1 and 2 of the definition of assume-admissible.

$\boxed{\mathsf{AA} \Rightarrow \mathsf{Coop}}$ This holds by Theorem 1.

$\boxed{\mathsf{Coop} \Rightarrow \mathsf{RS}^{\exists}(\mathsf{SPE})}$ Note that in order for RS to make sense we must have $\mathsf{sys} \in \mathcal{P}$. Assume Coop has a solution and let $\sigma_{\mathcal{P}}$ be a profile of strategy such that for all player i, $\sigma_{\mathcal{P}} \models \phi_i$.

We define a strategy profile σ'_i, that follows the path $\rho = \mathsf{Out}_{\mathsf{G}}(\sigma_i)$ when possible (that is: if h is a prefix of ρ then play $\mathsf{act}_{|h|}(\rho)$) and if not follows a subgame perfect equilibrium: that is, we select for each state s a subgame perfect equilibrium $\sigma^s_{\mathcal{P}}$, there always exist one for Borel games (so in particular for Muller games) [29, Theorem 3.15]; then if h is not a prefix of ρ, let j be the last index such that $h_{\leq j} = \rho_{\leq j}$ and we define $\sigma'_{\mathcal{P}}(h) = \sigma^{h_{j+1}}_{\mathcal{P}}(h_{\geq j+1})$.

Let h be a history. If h is a prefix of ρ then the objective of each player is satisfied by following $\sigma'_i \circ h$ so none of them can gain by changing its strategy, therefore it is a Nash equilibrium from $\mathsf{last}(h)$. If h is not a prefix of ρ then by definition of σ'_i, players follow a subgame-perfect equilibrium since h deviated from ρ, so in particular $\sigma'_i \circ h$ is a Nash equilibrium from $\mathsf{last}(h)$. Moreover the objective of the system is satisfied. Therefore $\sigma_{\mathcal{P}}$ is a solution to $\mathsf{RS}(\mathsf{SPE})$.

$\boxed{\mathsf{Win} \Rightarrow \mathsf{AG}^{\vee}}$ Let $\sigma_{\mathcal{P}}$ such that for each player i, σ_i is winning for ϕ_i. The first condition in the definition of $\mathsf{AG}^{\vee}$ is satisfied because for all player i, $\mathsf{Out}_{\mathsf{G}}(\sigma_{\mathcal{P}})$ satisfies ϕ_i. The second condition is satisfied because

for all strategy σ'_{-i}, we have that $\mathsf{Out}_{\mathsf{G}}(\sigma_i, \sigma'_{-i})$ satisfies ϕ_i, so in particular it satisfies $(\bigvee_{j \in \mathcal{P} \setminus \{i\}} \phi_j \Rightarrow \phi_i)$. Hence $\sigma_{\mathcal{P}}$ is a solution for $\mathsf{AG}^{\vee}$.

$\boxed{\mathsf{AG}^{\vee} \Rightarrow \mathsf{AG}^{\wedge}}$ This holds because the second condition in the definition of these rules is stronger for $\mathsf{AG}^{\vee}$.

$\boxed{\mathsf{AG}^{\wedge} \Rightarrow \mathsf{Coop}}$ This implication holds simply because of the condition 1 in the definition of assume-guarantee, which corresponds to the definition of Cooperative synthesis.

□

Lemma 9. *For all* $\gamma \in \{\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom}\}$, $\mathsf{Win} \Rightarrow \mathsf{RS}^{\forall}(\gamma)$, $\mathsf{RS}^{\forall}(\mathsf{Dom}) \not\Rightarrow \mathsf{RS}^{\forall}(\mathsf{NE})$ *and* $\mathsf{RS}^{\exists}(\mathsf{Dom}) \Rightarrow \mathsf{RS}^{\exists}(\mathsf{SPE})$.

Proof. $\boxed{\mathsf{Win} \Rightarrow \mathsf{RS}^{\forall}(\gamma)}$ Let $\sigma_{\mathcal{P}}$ be a strategy profile such that for each player i, σ_i is winning for ϕ_i.

We first show that $\Sigma^{\gamma}_{\mathsf{G},\sigma_1}$ is not empty. For $\gamma \in \{\mathsf{NE}, \mathsf{SPE}\}$ this is because there always exist a subgame perfect equilibrium for Borel games (so in particular for Muller games) [29, Theorem 3.15] and a subgame perfect equilibrium is a Nash equilibrium. For $\gamma = \mathsf{Dom}$, note that by definition of dominant strategies, winning strategies are dominant, so $\Sigma^{\mathsf{Dom}}_{\mathsf{G},\sigma_1}$ contains at least σ_{-1}.

Let σ'_{-1} be a strategy profile for $\mathcal{P} \setminus \{1\}$. Since σ_1 is a winning we have that $\mathsf{G}, \sigma_1, \sigma'_{-1} \models \phi_1$. Therefore σ_1 is a solution for $\mathsf{RS}^{\forall}(\gamma)$.

$\boxed{\mathsf{RS}^{\exists}(\mathsf{Dom}) \Rightarrow \mathsf{RS}^{\exists}(\mathsf{SPE})}$ Let $\sigma_{\mathcal{P}}$ be a witness for $\mathsf{RS}^{\exists}(\mathsf{Dom})$. We define a strategy profile $\sigma'_{\mathcal{P}}$ such that σ'_i follows σ_i on all histories compatible with σ_i (that is if h prefix of $\rho \in \mathsf{Out}_{\mathsf{G}}(\sigma_i)$ then $\sigma'_i(h) = \sigma_i(h)$) and outside of these histories follows a subgame perfect equilibria: there always exist one for Borel games (so in particular for Muller games) [29, Theorem 3.15].

By definition of $\sigma'_{\mathcal{P}}$, the outcome $\mathsf{Out}_{\mathsf{G}}(\sigma'_{\mathcal{P}})$ is the same than $\mathsf{Out}_{\mathsf{G}}(\sigma_{\mathcal{P}})$. Because $\sigma_{\mathcal{P}}$ is a witness for $\mathsf{RS}^{\forall}(\mathsf{Dom})$, this outcome is winning for player 1.

It remains to show that σ'_{-1} is a subgame perfect equilibria. Let h be a history, i be a player different from player 1, and σ''_i be a strategy for player i. We show that from h player i does not improve by switching from σ'_i to another strategy σ''_i, which will show that $\sigma'_{\mathcal{P}} \circ h$ is a Nash equilibrium from h.

If h is compatible with σ_i then σ'_i coincide with σ_i from this history, so $\mathsf{Out}_{\mathsf{G}}(\sigma'_i, \sigma_{-i}) = \mathsf{Out}_{\mathsf{G}}(\sigma_{\mathcal{P}})$. Since σ_i is a dominant strategy, if $\mathsf{G}, \sigma''_i, \sigma_{-i} \models \phi_i$ then $\mathsf{G}, \sigma_i, \sigma_{-i} \models \phi_i$ and therefore this implies that $\mathsf{Out}_{\mathsf{G}}(\sigma'_i, \sigma_{-i})$ satisfy ϕ_i. This means that i does not improve by switching from σ'_i to σ''_i.

If h is not compatible with σ_i, then σ'_i plays according to a subgame-perfect equilibria since the first deviation. In particular, this strategy is a Nash equilibrium from h.

This shows that σ'_{-1} is a subgame perfect equilibrium and has $\sigma'_{\mathcal{P}} \models \phi_1$, this is a witness for $\mathsf{RS}^{\exists}(\mathsf{SPE})$.

$\boxed{\mathsf{RS}^{\forall}(\mathsf{Dom}) \not\Rightarrow \mathsf{RS}^{\forall}(\mathsf{NE})}$ Consider the example given in Figure 4. The strategy r for player 2 is dominant and any strategy of player 3 is dominant. The outcome of these strategies always go to the bottom state where ϕ_{sys} is satisfied. Therefore there is a solution to $\mathsf{RS}^{\forall}(\mathsf{Dom})$.

However, we show that there is no solution to $\mathsf{RS}^{\forall}(\mathsf{NE})$. Consider the strategy profile $(\cdot, l, b)$, this is a Nash equilibrium (even a subgame Nash equilibrium) since no player can improve his/her strategy. Note that player 1 is losing for that profile, hence no strategy of player 1 can ensure that it will win for all Nash equilibria. □

In the example of Section 4, we saw that more strategy profiles satisfied the assume-guarantee condition compared to assume-admissibility, including undesirable strategy profiles. We show that the rule $\mathsf{AG}^{\wedge}$ is indeed more often satisfied than AA; while the rules $\mathsf{AG}^{\vee}$, and AA are incomparable.

Lemma 10. *We have* $\mathsf{AG}^{\wedge} \not\Rightarrow \mathsf{AA}$; $\mathsf{AG}^{\vee} \not\Rightarrow \mathsf{AA}$; $\mathsf{AA} \not\Rightarrow \mathsf{AG}^{\wedge}$ *and* $\mathsf{AA} \not\Rightarrow \mathsf{AG}^{\vee}$.

Proof. $\boxed{\mathsf{AG}^{\wedge} \not\Rightarrow \mathsf{AA} \text{ and } \mathsf{AG}^{\vee} \not\Rightarrow \mathsf{AA}}$ Consider the game represented in Figure 5. In this example, we have $\mathsf{Adm}_1 = \Sigma_1$. Therefore player 2 has no winning strategy against all admissible strategies of Adm_2 (in

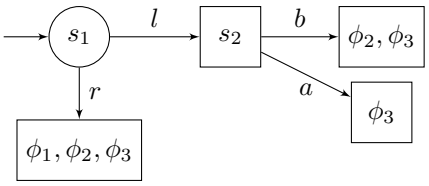

Figure 4: Example showing that $\mathsf{RS}^{\forall}(\mathsf{Dom}) \not\Rightarrow \mathsf{RS}^{\forall}(\mathsf{NE})$. Player 2 controls circle states, player 3 square states and player 1 does not control any state.

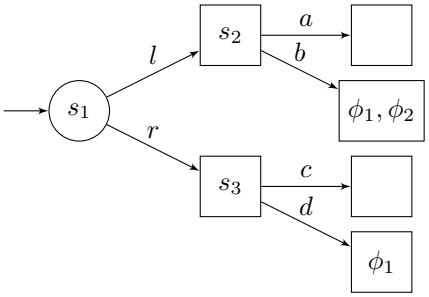

Figure 5: Example showing that $\mathsf{AG} \not\Rightarrow \mathsf{AA}$. Player 1 controls circle states and player 2 square states.

particular the strategy of player 1 that plays r, makes player 2 lose). So AA fails. However, we do have $\mathsf{AG}^{\wedge}$ by the profile $\sigma_1 \colon s_1 \mapsto l, \sigma_2 \colon s_2 \mapsto b, s_3 \mapsto c$. This profile also satisfies $\mathsf{AG}^{\vee}$ which is equivalent to $\mathsf{AG}^{\wedge}$ for two player games.

$\boxed{\mathsf{AA} \not\Rightarrow \mathsf{AG}^{\wedge}}$ Consider the example of Figure 6. The profile where player 1 and player 2 plays to the right is assume-admissible. However there is no solution to assume-guarantee synthesis: if player 1 and player 2 change their strategies to go to the state labeled ϕ_1, ϕ_2, then the condition $\mathcal{G}, \sigma_3 \models (\phi_1 \wedge \phi_2) \Rightarrow \phi_3$ is not satisfied.

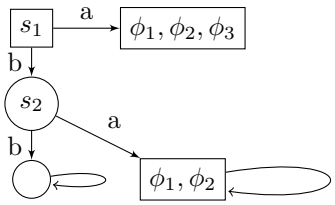

Figure 6: Example showing that $\mathsf{AA} \not\Rightarrow \mathsf{AG}^{\wedge}$. Player 1 controls circle states and player 2 square states; player 3 does not control any state.

$\boxed{\mathsf{AA} \not\Rightarrow \mathsf{AG}^{\vee}}$ We will provide a counter-example to show our claim. Note that we need strictly more than two players since otherwise $\mathsf{AG}^{\vee}$ is equivalent to $\mathsf{AG}^{\wedge}$, and we have just shown that AA implies $\mathsf{AG}^{\wedge}$.

Consider the game with three players in Fig. 7. Define the following objectives: $\phi_1 = \mathtt{GF}(s_4 \vee s_7)$, $\phi_2 = \mathtt{GF}(s_4 \vee s_6)$, $\phi_3 = \mathtt{true}$, where ϕ_i is player i's objective. These are actually reachability objectives since the game ends in absorbing states.

Now, action b is dominated at states s_2 and s_3 for player 2. Thus player 1 has a AA-winning strategy which consists in taking a at s_1. Player 2 has a winning strategy in the game (taking a at both states).

Player 3 has a AA-winning strategy too since actions b are eliminated for player 2. Therefore, there is an AA-winning strategy profile which ends in s_4.

On the other hand, there is no $\mathsf{AG}^\vee$ profile. In fact, player 1 has no winning strategy to ensure $\phi_2 \vee \phi_3 \Rightarrow \phi_1$, which is equivalent to ϕ_1 since $\phi_3 = \texttt{true}$.

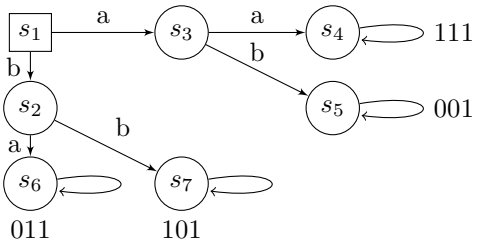

Figure 7: Example showing that $\mathsf{AA} \not\Rightarrow \mathsf{AG}^\vee$. Player 1 controls circle states and player 2 square states; player 3 does not control any state. At each absorbing state, the given Boolean vector represents the set of players for which the state is winning.

□

Lemma 11. *For two player games,* $\mathsf{AA} \Rightarrow \mathsf{AG}^\wedge$.

Proof. Assume $\mathcal{G}$ is a two player game and consider strategy profile (σ_1, σ_2) witness of AA. Note that if player j decreases its own value at position k then its value for $h_{\leq k+1}$ will be smaller or equal to 0 which means player j has no winning strategy from this history. By determinacy of turn-based zero-sum games, player $3-j$ has a winning strategy for $\neg\phi_j$. Therefore we can adjust the strategies (σ_1, σ_2) such that if there is a player j that decreases its own value, the other player will make it lose. We write (σ'_1, σ'_2) the strategies thus defined and we will show that they form a solution of Assume-Guarantee.

Let ρ be the outcome of the strategy profile (σ'_1, σ'_2). We can show that ρ is also the outcome of (σ_1, σ_2). First we recall that an admissible strategy does not decrease its own value (Lemma 2). Therefore each σ'_i is identical to σ_i on the run ρ. By Theorem 1, ρ satisfies $\phi_1 \wedge \phi_2$.

Let σ''_1 be an arbitrary strategy profile for 1, and consider $\rho' = \mathsf{Out}_\mathsf{G}(\sigma''_1, \sigma'_2)$. We show that $\rho' \models \phi_1 \Rightarrow \phi_2$. Note that player 2 cannot be the first to decrease its value during ρ' since it behave according to σ_2 has long has there are no devition, and σ_2 is admissible and admissible strategies do not decrease their own values.

- If player 1 decreases its value during ρ', player 2 will play to make him lose and $\rho' \not\models \phi_1$. As a consequence $\rho' \models \phi_1 \Rightarrow \phi_2$.
- Otherwise no player decreases its own value during ρ'. We assume that $\rho' \models \phi_1$ and show that $\rho' \models \phi_2$. Since $\rho' \models \phi_1$, by Lemma 13, there is a strategy τ''_1 which is admissible and compatible with ρ'. Since ρ' is an outcome of σ'_2, and of τ''_1, we have $\mathsf{Out}_\mathsf{G}(\tau''_1, \sigma_2) = \rho'$. Now, since τ''_1 is admissible and by the fact that σ_2 satisfies the condition 2 of AA, we obtain $\rho' \models \phi_2$, which proves the property.

We can show the same property replacing the roles of player 1 and player 2, thus showing that the profile is solution of $\mathsf{AG}^\wedge$. □

To prove that there is no implication for the edges that are not in the diagram of Figure 1, it remains to be shown that:

Lemma 12. $\mathsf{AA} \not\Rightarrow \mathsf{Win}$, $\mathsf{AG}^\vee \not\Rightarrow \mathsf{Win}$, $\mathsf{Coop} \not\Rightarrow \mathsf{AA}$, $\mathsf{Coop} \not\Rightarrow \mathsf{AG}^\wedge$, $\mathsf{Coop} \not\Rightarrow \mathsf{RS}^\exists(\mathsf{Dom})$, *and for all* $\gamma \in \{\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom}\}$, $\mathsf{RS}^{\exists,\forall}(\gamma) \not\Rightarrow \mathsf{Coop}$.

Proof. $\boxed{\mathsf{AA} \not\Rightarrow \mathsf{Win}}$ Towards a contradiction assume $\mathsf{AA} \Rightarrow \mathsf{Win}$, then since we have $\mathsf{Win} \Rightarrow \mathsf{AG}^\wedge$ (Lemma 8), we would have $\mathsf{AA} \Rightarrow \mathsf{AG}^\wedge$ but this contradicts Lemma 10.

$\boxed{\mathsf{AG}^\vee \not\Rightarrow \mathsf{Win}}$ Towards a contradiction assume $\mathsf{AG}^\vee \Rightarrow \mathsf{Win}$, then since we have $\mathsf{Win} \Rightarrow \mathsf{AA}$ (Lemma 8), we would have $\mathsf{AG}^\vee \Rightarrow \mathsf{AA}$ but this contradicts Lemma 10.

$\boxed{\mathsf{Coop} \not\Rightarrow \mathsf{AA}}$ In Figure 5, we have an example of a game where there is no solution for AA (see the proof of Lemma 10 for details), however there is a solution for Coop: (l, b).

$\boxed{\mathsf{Coop} \not\Rightarrow \mathsf{AG}^\wedge}$ Consider the example of Figure 8. There is a solution for Coop: player 1 plays a. However there is no solution for $\mathsf{AG}^\wedge$: player 2 has no strategy to ensure that $\phi_1 \implies \phi_2$.

$\boxed{\mathsf{Coop} \not\Rightarrow \mathsf{RS}^\exists(\mathsf{Dom})}$ Consider the example of Figure 9. This example has a solution for Coop, for instance (l, ac) or (r, bd). However player 2 has no dominant strategy: l looses against bd so it is dominated by r, and r looses against ac so it is dominated by l. Therefore $\mathsf{RS}^\exists(\mathsf{Dom})$ has no solution.

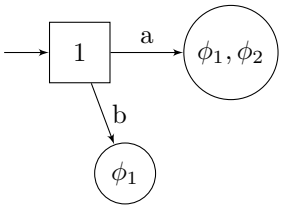

Figure 8: Example showing that $\mathsf{Coop} \not\Rightarrow \mathsf{AG}^\wedge$. Player 1 controls the circle state.

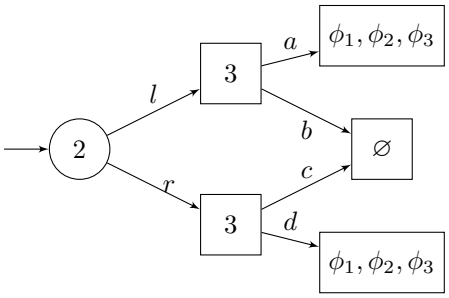

Figure 9: Example showing that $\mathsf{Coop} \not\Rightarrow \mathsf{RS}^\exists(\mathsf{Dom})$. Player 2 controls circle states, player 3 square states and player 1 does not control any state.

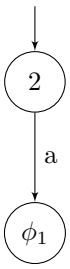

Figure 10: Two-player game showing that $\mathsf{RS}^{\exists,\forall}(\gamma) \not\Rightarrow \mathsf{Coop}$. Player 2 controls the circle state but has no choice.

$\boxed{\mathsf{RS}^{\exists,\forall}(\gamma) \not\Rightarrow \mathsf{Coop}}$ Consider the example of Figure 10. There is no solution for Coop: player 2 can never win. However there is a solution for any concept in $\mathsf{RS}^{\exists,\forall}(\gamma)$: player 1 wins against any of the strategy satisfying these concepts since the only possible outcome is winning for him.

□

This concludes the proof of Theorem 2.

B.2 Proof of Lemma 1

We now give the proof of Lemma 1

Lemma 1. *Let* $\mathsf{G} = \langle \mathsf{A}, \phi_1, \phi_2 \rangle$ *be a two-player game. If player* 2 *has a winning strategy for* ϕ_2 *and* Win-under-Hyp *has a solution, then* AA *has a solution.*

Proof. Assume that σ_2^w is a winning strategy for ϕ_2 and let σ_1, σ_2 be a solution of Win-under-Hyp. We have that $\forall \sigma_2'.\ \sigma_1, \sigma_2' \models \phi_2 \Rightarrow \phi_1$ and $\forall \sigma_1'.\ \sigma_1', \sigma_2 \models \phi_1 \Rightarrow \phi_2$. Since σ_2^w is a winning strategy, all admissible strategies of player 2 are winning. Then, for all $\sigma_2' \in \mathsf{Adm}_2$, we have $\mathsf{G}, \sigma_1, \sigma_2' \models \phi_2$ and because $\forall \sigma_2', \sigma_1, \sigma_2' \models \phi_2 \Rightarrow \phi_1$, we also have that $G, \sigma_1, \sigma_2' \models \phi_1$. If σ_1 is dominated, there exists a non-dominated strategy σ_1^a that dominates it [2, Thm. 11], otherwise we take $\sigma_1^a = \sigma_1$. In both cases σ_1^a is admissible. As σ_1 is dominated by σ_1^a, $\mathsf{G}, \sigma_1, \sigma_2' \models \phi_1$ implies $\mathsf{G}, \sigma_1^a, \sigma_2' \models \phi_1$. This shows that the condition $\forall \sigma_2' \in \mathsf{Adm}_2(\mathsf{G}).\ \mathsf{G}, \sigma_1^a, \sigma_2' \models \phi_1$ is satisfied. Since σ_2^w is winning, it is admissible and we also have $\forall \sigma_1' \in \mathsf{Adm}_1(\mathsf{G}).\ \mathsf{G}, \sigma_1', \sigma_2^w \models \phi_2$. Therefore all conditions of Assume-Admissible are satisfied by (σ_1^a, σ_2^w).

□

We now prove Theorem 3.

Theorem 3. *We have 1. Rule* AA *is rectangular; and for all games* G*,* AA*-winning strategy profile* σ_P*, coalition* $C \subseteq \mathcal{P}$*, if* $\sigma'_C \in \mathsf{Adm}_C(\mathsf{G})$*, then* $\mathsf{G}, \sigma_{-C}, \sigma'_C \models \bigwedge_{i \in -C} \phi_i$*. 2. The rules* Win *and* $\mathsf{RS}^\forall(\mathsf{Dom})$ *are rectangular; the rules* Coop*,* $\mathsf{AG}^\vee$*,* $\mathsf{AG}^\wedge$*,* $\mathsf{RS}^\exists(\mathsf{NE}, \mathsf{SPE}, \mathsf{Dom})$*, and* $\mathsf{RS}^\forall(\mathsf{NE}, \mathsf{SPE})$ *are not rectangular.*

Proof. $\boxed{\mathsf{AA} \text{ is rectangular}}$ If there is no solution to AA, then the set of witness is empty, and therefore is rectangular. If there is only one solution, then it is the Cartesian product of singletons and therefore also a rectangular set.

Otherwise let $\sigma_\mathcal{P}$ and $\sigma'_\mathcal{P}$ be two solutions of AA. Let i be a player of $\mathcal{P}$, we show that σ_i, σ'_{-i} is also a solution of AA. We have that $\sigma_i \in \mathsf{Adm}(\mathsf{G})$ and for all $j \neq i$, $\sigma_j \in \mathsf{Adm}(\mathsf{G})$, because condition 1 holds for $\sigma_\mathcal{P}$ and $\sigma'_\mathcal{P}$. Therefore condition 1 holds for σ_i, σ'_{-i}. Similarly, $\forall \sigma'_{-i} \in \mathsf{Adm}_{-i}(\mathsf{G})$. $\mathsf{G}, \sigma'_i, \sigma_i \models \phi_i$ and for all $j \neq i$, $\forall \sigma'_{-j} \in \mathsf{Adm}_{-j}(\mathsf{G})$. $\mathsf{G}, \sigma'_j, \sigma_j \models \phi_j$, because condition 2 holds for $\sigma_\mathcal{P}$ and $\sigma'_\mathcal{P}$. Therefore condition 2 holds for σ_i, σ'_{-i} and it is a witness of AA.

Let Σ_i^{aa} be the set of strategy σ_i such that there exists σ_{-i} such that σ_i, σ_{-i} is a witness of AA. We can show that the set of witness of AA is the Cartesian product of the Σ_i^{aa}. We obviously have that the set of solutions is included in $\prod_{i \in \mathcal{P}} \Sigma_i^{aa}$. Let $\sigma_\mathcal{P}$ be a profile in $\prod_{i \in \mathcal{P}} \Sigma_i^{aa}$, and $\sigma'_\mathcal{P}$ a witness of AA. We can replace for one i at a time, the strategy σ'_i by σ_i in $\sigma'_\mathcal{P}$ and by the small property we previously proved, the strategy profile stays a solution of AA. Therefore $\sigma_\mathcal{P}$ is a solution of AA. This shows that the set of solutions is the rectangular set $\prod_{i \in \mathcal{P}} \Sigma_i^{aa}$.

$\boxed{\sigma'_C \in \mathsf{Adm}_C(\mathsf{G}) \Rightarrow \mathsf{G}, \sigma_{-C}, \sigma'_C \models \bigwedge_{i \in -C} \phi_i}$ This claim follows from the definition of AA-winning strategy profiles, since each strategy is winning against admissible strategies.

Now Consider any game G and fix a profile $\sigma_\mathcal{P}$ such that $\mathsf{G}, \sigma_\mathcal{P} \models \bigwedge_{1 \leq i \leq n} \phi_i$.

$\boxed{\mathsf{Win}}$ Assume $\sigma_\mathcal{P}$ is solution to Win, then each σ_i is a winning strategy. Let σ'_i be a strategy part of another profile solution to Win. Then the strategy σ'_i ensures ϕ_i against any strategy profile for $-i$. If we replace σ_i by σ'_i in the profile $\sigma_\mathcal{P}$ then the condition for Win are still satisfied. Thus the rule Win is rectangular.

$\boxed{\mathsf{RS}^\forall(\mathsf{Dom})}$ Let σ_1 be a solution of $\mathsf{RS}^\forall(\mathsf{Dom})$. Let $\sigma_2, \ldots, \sigma_n$ and $\sigma'_2, \ldots, \sigma'_n$ be profiles of $\Sigma^{\mathsf{Dom}}_{\mathsf{G},\sigma_1}$ such that $\sigma_1, \sigma_2, \ldots, \sigma_n \models \phi_1$ and $\sigma_1, \sigma'_2, \ldots, \sigma'_n \models \phi_1$. If we define a profile $\tau_2, \ldots, \tau_n$ where each τ_i is either σ_i or σ'_i, then as each τ_i is dominant, we have $\sigma_1, \tau_2, \ldots, \tau_n \models \phi_1$ because σ_1 is a solution of $\mathsf{RS}^\forall(\mathsf{Dom})$. Therefore the profile belongs to $\Sigma^{\mathsf{Dom}}_{\mathsf{G},\sigma_1}$ and makes σ_1 win. This shows that the rule is rectangular.

$\boxed{\mathsf{RS}^\exists(\mathsf{Dom})}$ Consider the example of Figure 11. Since player 2 and player 3 are always winning, all their strategies are dominant. There is only one strategy σ_1 for player 1 since it controls no state. The profiles (a, c) and (b, d) are strategies of $\Sigma^{\mathsf{Dom}}_{\mathsf{G},\sigma_1}$ such that σ_1 wins for ϕ_1, but the profile (σ_1, a, d) does not make ϕ_1 hold. The rule is therefore not rectangular.

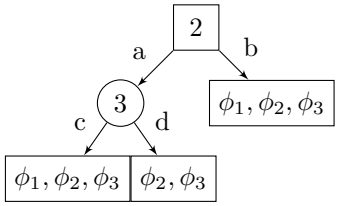

Figure 11: Game with three players showing that rule $\mathsf{RS}^\exists(\mathsf{Dom})$ is not rectangular. Here, player 1 controls no state; player 2 controls the square state, and player 3 controls the round state.

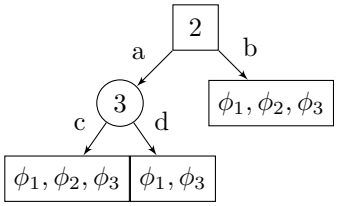

Figure 12: Game with three players showing that rule $\mathsf{RS}^\forall(\mathsf{NE})$ and $\mathsf{RS}^\forall(\mathsf{SPE})$ are not rectangular. Player 1 controls no state, player 2 controls the square state and player 3 the round state.

$\boxed{\mathsf{RS}^\forall(\mathsf{NE}) \text{ and } \mathsf{RS}^\forall(\mathsf{SPE})}$ Consider the game represented in Figure 12. Player 1 has only one strategy σ_1 and the other players have two possible strategies: a and b for player 2 and c and d for player 3. Since

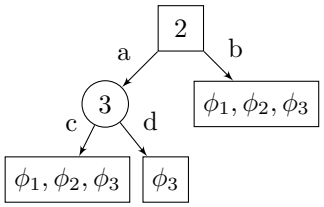

Figure 13: Game with three players showing that rule $\mathsf{RS}^{\exists}(\mathsf{NE})$ and $\mathsf{RS}^{\exists}(\mathsf{SPE})$ are not rectangular. Player 1 controls no state, player 2 controls the square state and player 3 the round state.

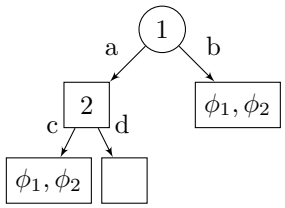

Figure 14: Game with two players showing that rule $\mathsf{AG}^{\vee}$ and $\mathsf{AG}^{\wedge}$ are not rectangular. Player 1 controls the round state and player 2 the square state.

player 1 is always winning, σ_1 is a solution for $\mathsf{RS}^{\forall}(\mathsf{NE}, \mathsf{SPE})$. The profiles (a, c) and (b, d) are two (subgame perfect) Nash equilibria which make ϕ_1 hold. However the profile (a, d) obtained by picking one strategy in each profile, is no longer a Nash equilibrium (and so not a subgame perfect equilibrium). Therefore $\mathsf{RS}^{\forall}(\mathsf{NE})$ and $\mathsf{RS}^{\forall}(\mathsf{SPE})$ are not rectangular.

$\boxed{\mathsf{RS}^{\exists}(\mathsf{NE}) \text{ and } \mathsf{RS}^{\exists}(\mathsf{SPE})}$ Consider the game represented in Figure 13. Player 1 has only one strategy and the other players have two possible strategies: a and b for player 2 and c and d for player 3. The profiles (a, c) and (b, d) are two (subgame perfect) Nash equilibria which make ϕ_1 hold. However the profile (a, d) obtained by taking one strategy in each profile, is no longer winning for player 1. Therefore $\mathsf{RS}^{\exists}(\mathsf{NE})$ and $\mathsf{RS}^{\exists}(\mathsf{SPE})$ are not rectangular.

$\boxed{\mathsf{Coop}}$ Once again, consider the game represented in Figure 13. The profiles (a, c) and (b, d) make all players win, but the profile (a, d), is no longer winning for the player 1, so it is not a solution of Coop. Therefore Coop is not rectangular.

$\boxed{\mathsf{AG}^{\vee} \text{ and } \mathsf{AG}^{\wedge}}$ Consider the game represented in Figure 14. The profiles (a, c) and (b, d) make the two players win. Since all possible outcome of the game satisfy the implications $\phi_1 \Rightarrow \phi_2$ and $\phi_2 \Rightarrow \phi_1$, both profiles are solution to $\mathsf{AG}^{\vee}$ and $\mathsf{AG}^{\wedge}$ (note that the two concepts coincide here because there are only two players). However the profile (a, d) obtained by taking one strategy in each profile, is no longer winning for the player 1. Therefore $\mathsf{AG}^{\vee}$ and $\mathsf{AG}^{\wedge}$ are not rectangular.

□

C Complements on Algorithm for Assume-Admissible Synthesis (Section 6)

In this section, we recall some detailed results on values and admissible strategies, and give the details of the algorithm for assume-admissible synthesis.

C.1 Values of Histories and Admissibility

For the proofs, we actually need a more refined notion of value, which coincides with the notion defined in the core of our paper for prefix-independent objectives. Let us give the formal definition from [2].

We fix a game G. Given a history h, and a set of strategies Σ'_i for player i, we write $\Sigma'_i(h)$ for the set of strategies of Σ'_i compatible with h, that is, the set of strategies σ_i such that $h \in \mathsf{Out}_{\mathsf{G}}(\sigma_i)$. We also write $\Sigma'(h)$ for $\prod_{i \in \mathcal{P}} \Sigma'_i(h)$.

Definition 1 (Value [2])**.** *Let Σ' be a rectangular set of strategy profiles. The* value of history h for player i with respect to Σ'*, written $\mathsf{Val}_i(\Sigma', h)$, is given by:*

- *if every $\sigma_{\mathcal{P}} \in \Sigma'(h)$ is losing for player i then $\mathsf{Val}_i(\Sigma', h) = -1$;*

- *if there is a strategy of* $\sigma_i \in \Sigma'_i(h)$ *such that for all strategy profiles* σ_{-i} *in* $\Sigma'_{-i}(h)$*, the profile* (σ_i, σ_{-i}) *is winning for player* i *then* $\mathsf{Val}_i(\Sigma', h) = 1$*;*
- *otherwise* $\mathsf{Val}_i(\Sigma', h) = 0$*;*

We use the shorthand notation $\mathsf{Val}_i(h) = \mathsf{Val}_i(\Sigma, h)$. Notice that the value only depends on the last state for prefix-independent objectives. We may thus write $\mathsf{Val}_i(s) = \mathsf{Val}_i(h)$ for $s = \mathsf{last}(h)$; observe that this is the notation we use in the core of the paper.

We noted in Lemma 2 that admissible strategies do not make the player's value decrease. We prove here that conversely, any winning run on which player i does not decrease its own value is compatible with an admissible strategy of player i.

Lemma 13. *Let* i *be a player, and* ρ *a history. If* $\rho \models \phi_i$ *and player* i *does not decrease its own value in any prefix of* ρ*, then there exists a strategy profile* $(\sigma_i, \sigma_{-i}) \in \mathsf{Adm}_i \times \Sigma_{-i}$ *such that* ρ *is the outcome of* (σ_i, σ_{-i})*.*

Proof. We define the strategies σ_i and σ_{-i} to follow ρ when possible (if the current history is a prefix of ρ, then proceed to the following state of ρ), and if a deviation has occurred in ρ, that is there is k such that $h_k = \rho_k$, $h_{k+1} \neq \rho_{k+1}$, then starting from $h_{\leq k+1}$, σ_i follows a non-dominated strategies with respect to $\Sigma(h_{k+1})$ (Thanks to [2] and the fact that this set allows shifting [2] such a non-dominated strategies exists). The run ρ is obviously an outcome of this profile. We now have to show that the strategy σ_i that we define is admissible. According to [2, Lem. 9], it is enough to show that for every history h outcome of σ_i, the value for player i with respect to $\{\sigma_i\} \times \Sigma_{-i}$ is greater or equal to that of Σ.

Let h be a finite outcome of σ_i. We distinguish the case where h has deviated from ρ and the case where it has not.

If a deviation has occurred, then σ_i follows a strategy non dominated with respect to $\Sigma(h_{\leq k+1})$ where k is the last index where $h_k = \rho_k$. By [2, Lem. 9], the value of $\{\sigma_i\} \times \Sigma_{-i}(h)$ in h is greater or equal to that of $\Sigma(h_{\leq k+1})$. Since $\Sigma_{-i}(h) \subseteq \Sigma_{-i}(h_{\leq k+1})$, the value of h with respect to $\{\sigma_i\} \times \Sigma_{-i}(h)$ is greater or equal to that with respect $\Sigma(h)$. Note that by the definition of the value, the value of h with respect to a rectangular set Σ' is equal to that of h with respect to $\Sigma'(h)$. Therefore the value of h with respect with $\{\sigma_i\} \times \Sigma$ is greater or equal to that with respect to Σ.

If a deviation has not occurred then h is a prefix of ρ. The value of h with respect to Σ is greater or equal to 0 since ρ is winning for ϕ_i. Then:

- if the value is 0, then as there is an outcome of σ_i after this history which is winning (the run ρ), the value of σ_i is at least 0;
- if the value is 1, then we can show that from this state σ_i plays a winning strategy: if we stay along ρ, the run is winning, if we deviate, the run goes to a state of value 1 (player i has not decreased its own value and from a state of value 1 for player i, any action of the adversaries lead to a state of value 1), and therefore σ_i revert to a winning strategy.

Therefore the property is satisfied by σ_i and it is admissible.

□

We now prove Lemma 3 characterizing the outcomes of admissible strategies. This result follows from [4], but we adapt it to make the formulas Φ_i appear explicitly.

Lemma 3. *For all games* G*, and players* i*,* $\mathsf{Out}_\mathsf{G} \cap \Phi_i = \mathsf{Out}_\mathsf{G}(\mathsf{Adm}_i, \Sigma_{-i})$*, where* $\Phi_i = \mathtt{G}(E_i) \wedge (\mathtt{GF}(V_{i,1}) \Rightarrow \phi_i) \wedge (\mathtt{GF}(V_{i,0}) \Rightarrow \phi_i \vee \mathtt{GF}(H_i))$*.*

Proof. In [4, Lemma 6], an automaton $\mathcal{A}^1_i$ is defined such that $\mathcal{A}^1_i \cap \mathsf{Out}_\mathsf{G}(\Sigma) = \mathsf{Out}_\mathsf{G}(\mathsf{Adm}_i, \Sigma_{-i})$. (Note that a more general construction $\mathcal{A}^n_i$ was given in [4]; we only need the case $n = 1$ here.)

We now analyze further the language of $\mathcal{A}^1_i$. The edges are those of $\mathtt{G}$ except for edges outside of E_i (these edges are noted T in [4]), so the set of runs in $\mathcal{A}^1_i$ corresponds to $\mathsf{Out}_\mathsf{G} \cap \mathtt{G}(E_i)$. Now a run of $\mathcal{A}^1_i$ is accepted if, and only if one the following condition is satisfied, writing $\mathsf{VR}(\rho)$ for the sequence $(\mathsf{Val}(\rho_i))_{i \in \mathbb{N}}$:

- $\mathsf{VR}(\rho) \in 0^*(-1)^\omega$;
- $\mathsf{VR}(\rho) \in 0^*1^\omega$ and $\rho \models \phi_i$;
- $\mathsf{VR}(\rho) \in 0^\omega$ and $\rho \models \phi_i$ or $\rho \models \mathtt{GF}(H_i)$.

Any run of $\mathsf{Out}_\mathsf{G} \cap \mathtt{G}(E_i)$ reaching some state of value -1 is necessarily losing; thus all successors also have value -1. Similarly, because we removed edges where player i decreases its own value, once the runs reaches a state of value 1, it never gets out of these states. Therefore runs of $\mathsf{Out}_\mathsf{G} \cap \mathtt{G}(E_i)$ have one of the three forms: $0^*(-1)^\omega$, 0^*1^ω or 0^ω.

Let ρ be a run that is accepted by $\mathcal{A}_i^1$, it satisfies $\mathtt{G}(E_i)$ and:

- if ρ ends in the states of value -1 then it does not visit $V_{i,1}$ or $V_{i,0}$ infinitely often and thus belongs to Φ_i;
- if ρ ends in the states of value 1, then by the acceptance condition it satisfies ϕ_i and thus belongs to Φ_i;
- otherwise it stays in the states of value 0, then by the acceptance condition, either it satisfies ϕ_i or $\mathtt{GF}(H_i)$ and thus belongs to Φ_i.

Now let ρ be a run that satisfies ϕ_i, it satisfies $\mathtt{G}(E_i)$ and therefore corresponds to a valid run of $\mathcal{A}_i^1$.

- If ρ ends in the states of value -1 then condition $\mathsf{VR}(\rho) \in 0^*(-1)^\omega$ is satisfied, thus ρ is accepted by $\mathcal{A}_i^1$.
- If ρ ends in the states of value 1, then by definition of Φ_i it satisfies ϕ_i and condition $\mathsf{VR}(\rho) \in 0^*1^\omega$ and $\rho \models \phi_i$ is satisfied, thus ρ is accepted by $\mathcal{A}_i^1$.
- Otherwise it stays in the states of value 0, then by definition of Φ_i, either ϕ_i or $\mathtt{GF}(H_i)$ holds for ρ, hence $\mathsf{VR}(\rho) \in 0^\omega$ and $\rho \models \phi_i$ or $\rho \models \mathtt{GF}(H_i)$ is satisfied, thus ρ is accepted by $\mathcal{A}_i^1$.

This shows that $\Phi_1 \cap \mathsf{Out}_\mathsf{G} = \mathcal{A}_i^1 \cap \mathsf{Out}_\mathsf{G}$ and by [4, Lemma 6], this equals $\mathsf{Out}(\mathsf{Adm}_i, \Sigma_{-i})$.

□

C.2 Algorithm

We start by a characterization of games with AA-winning strategy profiles. An assume-admissible winning strategy ensures its objective against all admissible strategies of other players; we formalize this in the following lemma using values (see Definition 1).

Lemma 14. *For all game* G*, there is an* AA*-winning strategy profile if, and only if for all players* i*,* $\mathsf{Val}_i((\mathsf{Adm}_j(\mathsf{G}))_{j\in\mathcal{P}}, s_\mathsf{init}) = 1$.

Proof. By definition of the value we have for all player i, that:

$$\mathsf{Val}_i((\mathsf{Adm}_i(\mathsf{G}))_{i\in\mathcal{P}}, s_\mathsf{init}) = 1 \quad \Leftrightarrow \quad \exists \sigma_i \in \mathsf{Adm}_i(\mathsf{G}).\ \forall \sigma_{-i} \in \mathsf{Adm}_{-i}(\mathsf{G}).\ \mathsf{G}, \sigma_i, \sigma_{-i} \models \phi_i \tag{1}$$

$\boxed{\Rightarrow}$ If $\sigma_\mathcal{P}$ is a witness for Assume-Admissible then for all players i, $\sigma_i \in \mathsf{Adm}_i(\mathsf{G})$ is such that for all $\sigma_{-i} \in \mathsf{Adm}_{-i}(\mathsf{G})$, $\mathsf{G}, \sigma_i, \sigma_{-i} \models \phi_i$. Hence by (1), for all players i, $\mathsf{Val}_i((\mathsf{Adm}_i(\mathsf{G}))_{i\in\mathcal{P}}, s_\mathsf{init}) = 1$.

$\boxed{\Leftarrow}$ Assume that for all players i, $\mathsf{Val}_i((\mathsf{Adm}_i(\mathsf{G}))_{i\in\mathcal{P}}, s_\mathsf{init}) = 1$. By (1) we have that for all players i, there exist $\sigma_i \in \mathsf{Adm}_i(\mathsf{G})$ such that $\forall \sigma_{-i} \in \mathsf{Adm}_{-i}(\mathsf{G}).\ \mathsf{G}, \sigma_i, \sigma_{-i} \models \phi_i$. Let $\sigma_\mathcal{P}$ be a strategy profile made of such strategies σ_i. We have that for all player i, $\sigma_i \in \mathsf{Adm}_i(\mathsf{G})$ (condition 1 in the definition of Assume-Admissible) and moreover for all $\sigma'_{-i} \in \mathsf{Adm}_{-i}(\mathsf{G})$, $\mathsf{G}, \sigma_i, \sigma'_{-i} \models \phi_i$ (condition 2 in the definition of Assume-Admissible). Hence the strategy profile $\sigma_\mathcal{P}$ constitutes a solution to Assume-Admissible.

□

The algorithm to synthesize AA-winning strategies uses procedures from [4], originally developed to compute the *outcomes* that survive the *iterative elimination* of dominated strategies. More precisely, the elimination procedure of [4] first computes the outcomes of admissible strategies; from this it deduces the strategies that are not dominated when all players are restricted to admissible strategies, and their possible outcomes; and this is repeated until the set of outcomes stabilizes. In the end, one obtains the set of the runs that are the outcomes of strategy profiles that have survived this iterative elimination.

Here, we roughly consider the first iteration of the above procedure, and explicitly give algorithms to actually synthesize strategies that are winning against admissible strategies.

We explain the proof of Lemma 4 which follows from [4].

Lemma 4. *A player i strategy is AA-winning iff it is winning for objective Ω_i.*

Proof. It is shown in [4, Prop. 5] that a strategy of player i is a strategy of Σ_i^n which is winning from state s against all strategies of Σ_{-i}^n if, and only if, it is winning for objective $\Omega_i^n(s)$ (where Σ^n is the set of strategy that remain after n step of elimination). The results immediately follows of the case $n = 1$.

□

We will now establish the correctness of the game G'_i. Let us first formalize the correspondence between G and G'_i. We define relation $\sim \subseteq \mathsf{S} \times \mathsf{S}'$: for all $s \in \mathsf{S} \times \{\bot, 0, \top\}$, $s \sim (s, x)$. We extend this to runs by $\rho \sim \rho'$ iff for all $i \in \mathbb{N}$, $\rho_i \sim \rho'_i$. In fact the relation defines a bijection:

Lemma 15. *For any $\rho \in \mathsf{Out}_{\mathsf{G}}$ there is a unique $\rho' \in \mathsf{Out}_{\mathsf{G}'_i}$ such that $\rho \sim \rho'$.*

Proof. Assume towards a contradiction that we have ρ' and ρ'' such that $\rho = \pi(\rho') = \pi(\rho'')$. Let i be the last state such that they coincide: $\rho'_i = \rho''_i$ and $\rho'_{i+1} \neq \rho''_{i+1}$. Since $\pi(\rho') = \pi(\rho'')$ we have that they differ only by the second component. We can assume without loss of generality that there is are actions a and b such that $(\rho_i, a) \in E_j$ (where player j controls ρ_i), $(\rho_i, b) \notin E_j$ and $\delta(\rho_i, a) = \rho_{i+1} = \delta(\rho_i, b)$. This means that there is are actions a and b such that $(s, a) \in E_j$ (where player j controls ρ_i), $(s, b) \notin E_j$ and $\delta(\rho_i, a) = \rho_{i+1} = \delta(\rho_i, b)$. We have $\delta(\rho_i, b) = \delta(\rho_i, a)$, then by definition of E_j and because $(\rho_i, a) \in E_j$, $\mathsf{Val}_j(\delta(s, a)) = \mathsf{Val}_j(s)$ therefore $\mathsf{Val}_j(\delta(s, b)) = \mathsf{Val}_j(s)$ and by definition of E_j, (ρ_i, b) belongs to E_j which contradicts our assumptions and ends the proof.

□

We write π the bijection which, to $\rho' \in \mathsf{Out}_{\mathsf{G}'_i}$ associates $\rho \in \mathsf{Out}_{\mathsf{G}}$ with $\rho \sim \rho'$. We extend π as a mapping from strategies of G'_i to strategies of G by $\pi(\sigma'_i)(h) = \sigma'_i(\pi^{-1}(h))$. We can notice that for all strategies σ'_i, $\pi(\mathsf{Out}_{\mathsf{G}'_i}(\sigma'_i)) = \mathsf{Out}_{\mathsf{G}}(\pi(\sigma'_i))$.

Lemma 16. *Let G be a game, and i a player. Player i has a winning strategy for Ω_i in G if, and only if, he has a winning strategy for Ω'_i in G'_i. Moreover if σ'_i is winning for Ω'_i in G'_i then $\pi(\sigma'_i)$ is winning for Ω_i in G.*

Proof. We will first rewrite Ω_i in a form that is closer to that of Ω'_i. The objective Ω_i is defined by

$\mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_i) \cap (\mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_{-i}) \Rightarrow \phi_i)$. Observe that $\mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_{-i}) = \cap_{j \neq i} \mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_j)$ by definition.

$$\Omega_i = \mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_i) \cap \left(\bigcap_{j \neq i} \mathsf{Out}_{\mathsf{G}}(\mathsf{Adm}_j) \Rightarrow \phi_i \right)$$

$$\Omega_i = \Phi_i \cap \mathsf{Out}_{\mathsf{G}} \cap \left(\left(\mathsf{Out}_{\mathsf{G}} \cap \bigcap_{j \neq i} \Phi_j \right) \Rightarrow \phi_i \right) \text{ using Lem. 3}$$

$$\Omega_i = \Phi_i \cap \mathsf{Out}_{\mathsf{G}} \cap \left(\bigcap_{j \neq i} \Phi_j \Rightarrow \phi_i \right)$$

$$\Omega_i = \mathsf{Out}_{\mathsf{G}} \wedge \mathtt{G}(E_i) \wedge M_i \wedge \left(\left(\bigcap_{j \neq i} M_j \Rightarrow \phi_i \right) \vee \bigvee_{j \neq i} \neg \mathtt{G}(E_j) \right)$$

$$\Omega_i = \mathsf{Out}_{\mathsf{G}} \wedge \mathtt{G}(E_i) \wedge M_i \wedge \left(\left(\bigcap_{j \neq i} M_j \Rightarrow \phi_i \right) \vee \bigvee_{j \neq i} \mathtt{F}(\neg E_j) \right)$$

$\boxed{\Rightarrow}$ Let σ_i be a winning strategy for Ω_i in G. We consider the strategy σ'_i defined by $\sigma'_i(h') = \sigma_i(\pi(h'))$ and will show that it is winning for Ω'_i. Let ρ' be an outcome of σ'_i. We have that $\pi(\rho')$ is an outcome of σ_i. Since σ_i is winning for Ω_i, $\pi(\rho')$ belongs to Ω_i.

- If $\pi(\rho') \models M_i \wedge \mathtt{G}(E_i) \wedge \bigvee_{j \neq i} \mathtt{F}(\neg E_j)$, then by construction of δ' the play ρ' reaches a state of $\mathsf{S} \times \{\top\}$ and, from there, only states of $\mathsf{S} \times \{\top\}$ are visited. The condition $\mathtt{GF}(\mathsf{S} \times \{\top\}) \wedge M_i$ is met by ρ' and therefore ρ' is winning for Ω'_i.
- Otherwise $\pi(\rho') \models M_i \wedge \mathtt{G}(E_i) \wedge (\bigwedge_{j \neq i} M_j) \Rightarrow \phi_i$. By construction of δ' the play ρ' stays in $\mathsf{S} \times \{0\}$. The condition $\mathtt{GF}(\mathsf{S} \times \{0\}) \wedge M_i \wedge (\bigwedge_{j \neq i} M_j) \Rightarrow \phi_i$ is met by ρ' and therefore $\rho' \in \Omega'_i$.

This shows that the strategy σ'_i is winning for Ω'_i in G'_i.

$\boxed{\Leftarrow}$ Let σ'_i be a winning strategy for Ω'_i in G'_i, we show that $\pi(\sigma'_i)$ is winning for Ω_i in G. Let ρ be an outcome of $\pi(\sigma'_i)$. We have that $\pi^{-1}(\rho)$ is an outcome of σ'_i. Since σ'_i is winning for Ω'_i, $\pi^{-1}(\rho)$ belongs to Ω'_i. We have that $\pi^{-1}(\rho) \models \mathtt{GF}(\mathsf{S} \times \{0, \top\})$ and by construction of δ' this ensures that all edges that are taken belong to E_i and thus $\pi^{-1}(\rho)$ satisfies the condition $\mathtt{G}(E_i)$.

- If $\pi^{-1}(\rho) \models \mathtt{GF}(S \times \{\top\}) \wedge M_i$ then by construction of δ', an edge outside of E_j for some $j \neq i$ is taken. This ensures condition $\mathtt{F}(\neg E_j)$ and therefore ρ belongs to Ω_i.
- otherwise $\pi^{-1}(\rho) \models (\bigwedge_{j \neq i} M_j \Rightarrow \phi_i)$ and therefore ρ satisfies the condition $\mathtt{G}(E_i) \wedge M_i \wedge \left(\bigcap_{j \neq i} M_j \Rightarrow \phi_i \right)$ and hence belongs to Ω_i.

This shows that the strategy $\pi(\sigma'_i)$ is winning for Ω_i in G.

□

This characterization immediately yields the decidability of the problem in PSPACE. Moreover, PSPACE-hardness follows from that of Muller games. We also provide a polynomial-time algorithms for the particular case of Büchi conditions. The following theorem proves the first statement of Theorem 4; the second statement is proved afterwards (Theorem 7).

Theorem 6. AA*-synthesis in multiplayer Muller games is* PSPACE*-complete and* P*-complete for Büchi objectives.*

Proof. We use the result of [4, Prop. 8] in the special case where $n = 1$, in which case the proposition says that for multi-player games with Muller objectives, checking whether $\mathsf{Val}_i((\mathsf{Adm}_i(\mathsf{G}))_{i\in\mathcal{P}}, s_{\mathsf{init}}) = 1$ is in PSPACE. By Lem. 14, with one call to this algorithm for each player, we are able to decide if AA has a solution and this proves PSPACE membership. For hardness we encode 2-players zero-sum Muller games in our setting: the first player keep the same objective and the second one is always winning. In this setting no strategy of the second player is dominated so finding an AA winning strategy for the first player is the same as finding a winning strategy in the original game.

Büchi Case In the case of Büchi objectives, let us write $\phi_i = \mathtt{GF}(B_i)$ and consider the objective $M_i = \mathtt{GF}(V_{i,1}) \Rightarrow \mathtt{GF}(B_i) \wedge \mathtt{GF}(V_{i,0}) \Rightarrow (\mathtt{GF}(B_i) \vee \mathtt{GF}(H_i))$. In G, a run that verifies $\mathtt{G}(E_i)$ will either visit only $V_{i,1}$ after some point, or only $V_{i,-1}$ after some point, or only $V_{i,0}$ (see the proof of Lemma 3 for details). This means that M_i coincide with $\mathtt{GF}((V_{i,1} \wedge B_i) \vee (V_{i,0} \wedge B_i) \vee (V_{i,0} \wedge H_i) \vee V_{i,-1})$ on the language $\mathsf{Out}_{\mathsf{G}}(\Sigma) \cap \mathtt{G}(E_i)$. This is a Büchi condition. Note that in G'_i, runs that take an edge outside of E_i end in the $\mathsf{S} \times \{\bot\}$ component and Ω'_i disallows such runs. Therefore, in the expression of Ω'_i, we can also replace M'_i by a Büchi condition that we write $\mathtt{GF}(B_i^M)$. By abuse of notation we will also write B_i for the states $B_i \times \{\bot, 0, \top\}$ of the game G'_i. Consider now the objective Ω'_i, it is given by $(\mathtt{GF}(\mathsf{S} \times \{0\}) \wedge \mathtt{GF}(B_i^M) \wedge (\bigwedge_{j\neq i} \mathtt{GF}(B_j^M) \Rightarrow \mathtt{GF}(B_i \times \{\bot, 0, \top\}))) \vee (\mathtt{GF}(\mathsf{S} \times \{\top\}) \wedge \mathtt{GF}(B_i^M))$. Since in this game, states of $\mathsf{S} \times \top$ and $\mathsf{S} \times \bot$ are absorbing (no play can get out of those components) we write an equivalent objective which is: $(\mathtt{GF}(B_i^M \times \{0\}) \wedge (\bigwedge_{j\neq i} \mathtt{GF}(B_j^M) \Rightarrow \mathtt{GF}(B_i))) \vee (\mathtt{GF}(B_i^M \times \{\top\})$. We define a (small) deterministic parity automaton $\mathcal{A}$ that recognizes this language. Its state space is $(\{s, t, u, v\} \times (\{j \mid j \in \mathcal{P} \setminus \{i\}\} \cup \{\top\}))$, the transition relation is a product of transitions for the two components: $s \xrightarrow{B_i^M \times \{0,\top\}} u$, $u \xrightarrow{\neg B_i^M} t$, $t, u \xrightarrow{B_i^M \setminus B_i}$, $t, u \xrightarrow{B_i} v$, $v \to s$, and $j \xrightarrow{\neg B_j \times \{0\}} j$, $j \xrightarrow{B_j \times \{0\}} j'$ where j' is $j + 1$ if $j + 1 \in \mathcal{P} \setminus \{i\}$, $j + 2$ if $j + 1 = i$ and $j + 2 \in \mathcal{P}$, $\top$ otherwise, $\top \to j_0$ where j_0 is the smallest element of $\mathcal{P} \setminus \{i\}$. The coloring is defined by a function χ where $\chi(v, *) = 4$ (where $*$ is any possible second component), $\chi(\{s, t, u\}, \top) = 3)$, $\chi(u, \mathcal{P} \setminus \{i\}) = 2$, and for all other states s, $\chi(s) = 1$. A word is accepted by $\mathcal{A}$ when the maximal color appearing infinitely often is even.

We show that a play of G'_i satisfies Ω'_i if, and only if, it is a word accepted by $\mathcal{A}$.

Let ρ be a play of G'_i which satisfies Ω'_i, either it ends in the $\mathsf{S} \times \top$ component or the $\mathsf{S} \times 0$ component:

- If ρ ends in the $\top$ component then the state of color 3 will not be visited infinitely often (we need to be in $\top$ states to progress on this component of the automaton). As ρ visits infinitely often B_i^M, the corresponding run in $\mathcal{A}$ will visit infinitely often u, and therefore the maximal color that appears infinitely often is at least 2; and since it is not 3, it has to be even.
- Otherwise ρ ends in the 0 component. Since ρ satisfies Ω'_i, it visits B_i^M infinitely often and either there is a B_j^M for $j \neq i$ that is not visited infinitely often, or ρ visits infinitely often B_i.
  - If there is a B_j^M for $j \neq i$ that is not visited infinitely often, then the second component of $\mathcal{A}$ will get stuck at some point and its state $\top$ will not be visited infinitely often the state of color 3. As ρ visits infinitely often B_i^M, the corresponding run in $\mathcal{A}$ will visit infinitely often u, and therefore the maximal color that appears infinitely often is at least 2; and since it is not 3, it has to be even.
  - Otherwise ρ visits infinitely often B_i. Since we also visit B_i^M infinitely often, the run of $\mathcal{A}$ corresponding to ρ will reach infinitely often a state $(v, *)$ and therefore the maximal color occurring infinitely often is 4.

This proves that the word is accepted by $\mathcal{A}$.

Now let ρ be a play of G'_i such that the corresponding word is accepted by $\mathcal{A}$. If it is accepted then either the color 4 is seen infinitely often or the color 2 is and the color 3 is not:

- If the color 4 is visited infinitely often then this means t is reached infinitely often, and because of the structure of $\mathcal{A}$, u also is, which means both $B_i^M \times \{0\}$ and B_i occur infinitely often. This implies that the run ρ belongs to Ω'_i.

- Otherwise the color 2 is visited infinitely often and 3 is not. The states $(*, \top)$ are therefore not visited infinitely often (otherwise the maximal color would be 3 or 4. We deduce from that and the structure of $\mathcal{A}$ that some B_j^M for $j \neq i$ is not visited infinitely often. This means $\bigwedge_{j \neq i} \mathtt{GF}(B_j^M)$ is not true for ρ. Since the color 2 is seen infinitely often, this means $u, *$ is seen infinitely often and therefore $B_i \times \{0, \top\}$. This ensures ρ belongs to Ω_i'.

This proves that a play of G_i' satisfy Ω_i' if, and only if, it is a word accepted by $\mathcal{A}$.

Then solving the game G_i' with objective Ω_i' is the same as solving it with objective given by $\mathcal{A}$. This can be done by solving the parity game obtained by the product of G_i' with the automaton $\mathcal{A}$. The obtained game is of polynomial size and the number of priority is 4, such games can be solved in polynomial time (see for instance [22, 28]) and therefore we can decide our problem in polynomial time.

□

We are now interested in *computing* an AA-winning strategy profile. Thanks to Lemma 16, we obtain an algorithm to compute AA-winning strategies by looking for winning strategies in G_i' and projecting them:

Theorem 7. *Given a game* G *with Muller objectives, if* AA *has a solution, then an* AA*-winning strategy profile can be computed in exponential time.*

Proof. If AA has a solution, then by Lemma 16, there is a winning strategy for Ω_i' in G_i'. This Muller game has polynomial size, hence we can compute a winning strategy σ_i' in exponential time (for instance in [26] the authors show that we can compute such a winning strategy via a safety game of size $|\mathsf{S}|!^3$). By Lemma 16, the projection $\pi(\sigma_i')$ is an AA-winning strategy. Doing this for each player we obtain a strategy profile solution of AA.

□

D Complements on Abstraction (Section 7)

D.1 Abstract Games

The pair of abstraction and concretization functions (α, γ) actually defines a *Galois connection*:

Lemma 17. *The pair* (α, γ) *is a* Galois connection*, that is, for all* $S \subseteq \mathsf{S}$ *and* $T \subseteq \mathsf{S}^{\mathsf{a}}$*, we have that* $\alpha(S) \subseteq T$ *if, and only if,* $S \subseteq \gamma(T)$*.*

Proof. $\boxed{\Rightarrow}$ Let $s \in S$. Since γ defines a partition of S, there exists $t \in \mathsf{S}^{\mathsf{a}}$ such that $s \in \gamma(t)$. By definition of α, $\alpha(s) = t$. Assuming $\alpha(S) \subseteq T$, we have that $t \in T$. As $s \in \gamma(t)$, we have $s \in \gamma(T)$.

$\boxed{\Leftarrow}$ If $s^{\mathsf{a}} \in \alpha(S)$, then there is $s \in S$ such that $s^{\mathsf{a}} = \alpha(s)$. Assuming $S \subseteq \gamma(T)$, there is $t \in T$ such that $s = \gamma(t)$. By definition of α, we have that $\alpha(s) = t$. Therefore $s^{\mathsf{a}} \in T$.

□

We prove the soundness of the abstract arenas $\mathcal{A}^C$ we defined, by showing that if coalition C achieves an objective in G, then it also achieves the objective in $\mathcal{A}^C$.

Lemma 5. *If the coalition* C *has a winning strategy for objective* ϕ_k *in* G *from* s *then it has a winning strategy for* ϕ_k^{a} *in* $\mathcal{A}^C$ *from* $\alpha(s)$*.*

Proof. Assume σ_C is a winning profile of coalition C, for objective ϕ_k in G. We define by induction a winning strategy σ_C^{a} in $\mathsf{G}^{\mathsf{a},k,C}$. We assume that σ_C^{a} has been defined in a manner such that for each finite outcome h^{a} of σ_C^{a} shorter than some bound m, there is some $h \in \gamma(h^{\mathsf{a}})$ such that h is a finite outcome of σ_C. The idea is then to define σ_C^{a} to resolve the determinism in a way which simulates the behavior from h.

- If $s_i^{\mathsf{a}} \in \bigcup_{i \in \mathcal{P}} \mathsf{S}_i^{\mathsf{a}} \times \mathsf{Act}_i$, then $\sigma_C^{\mathsf{a}}(h^{\mathsf{a}} \cdot (\mathsf{last}(h^{\mathsf{a}}), a)) = \gamma(t)$ where $t = \delta(\mathsf{last}(h), a)$.
- If $s^{\mathsf{a}} \in \bigcup_{i \in C} \mathsf{S}_i^{\mathsf{a}}$, $\sigma_C^{\mathsf{a}}(h^{\mathsf{a}} \cdot (\mathsf{last}(h^{\mathsf{a}}), a) \cdot s^{\mathsf{a}}) = \sigma_C(h \cdot \delta(\mathsf{last}(h), a))$.

With this definition, our induction hypothesis will be respected for histories containing one more step, and therefore this holds for all histories. Let now ρ^{a} be an outcome σ_C^{a}. By the way we defined this strategy there is a run ρ outcome of σ_C such that $\rho \in \gamma(\rho^{\mathsf{a}})$. As σ_C is winning, ρ satisfies the Muller condition ϕ_k and since γ is compatible with players' objectives, ρ^{a} satisfies ϕ_k^{a}. Which show that C has a winning strategy in $\mathsf{G}^{\mathsf{a},k,C}$ for ϕ_k^{a}.

□

D.2 Abstract Assume-Admissible Synthesis

Value-Preserving Strategies We provide the proofs of the lemmas stated in the core of the paper.

Lemma 6. *For all players k and $x \in \{-1, 0, 1\}$, $\gamma(\underline{V}_{k,x}) \subseteq V_{k,x} \subseteq \gamma(\overline{V}_{k,x})$.*

Proof. $\boxed{\overline{V}_{k,1}}$ This is a direct consequence of Lemma 5.

$\boxed{\overline{V}_{k,-1}}$ If $s \in V_{k,-1}$ then the coalition $\mathcal{P}$ has no winning strategy in G. By determinacy, the empty coalition has a strategy to ensure $\neg\phi_k$. Therefore by Lemma 5, the coalition $\varnothing$ has a strategy in $\mathcal{A}^{\mathcal{P}}$ from $\alpha(s)$ that ensures $\neg\phi_k$. Therefore $s \in \gamma(\overline{V}_{k,-1})$.

$\boxed{\overline{V}_{k,0}}$ Recall that $V_{k,0} = \mathsf{Win}_{\mathcal{P}\setminus\{k\}}(\mathsf{A}, \neg\phi_k) \cap \mathsf{Win}_{\mathcal{P}}(\mathsf{A}, \phi_k)$. Let s be a state in $V_{k,0}$. By Lemma 5, $\alpha(s)$ belongs to both sides of the intersection, thus $\alpha(s) \in \overline{V}_{k,0}$. Thus $V_{k,0} \subseteq \gamma(\overline{V}_{k,0})$.

$\boxed{\underline{V}_{k,1}}$ If $s^{\mathsf{a}} \in \underline{V}_{k,1}$ then the coalition $\mathcal{P} \setminus \{k\}$ has no strategy in $\mathcal{A}^{\mathcal{P}\setminus\{k\}}$ for $\neg\phi_k^{\mathsf{a}}$. Therefore by Lemma 5, it has no strategy in A from any state of $\gamma(s^{\mathsf{a}})$ to do so. Therefore k has a winning strategy in A from $\gamma(s^{\mathsf{a}})$, and $\gamma(s^{\mathsf{a}}) \in V_{k,1}$.

$\boxed{\underline{V}_{k,-1}}$ If $s^{\mathsf{a}} \in \underline{V}_{k,-1}$, then the coalition $\mathcal{P}$ has no winning strategy in $\mathcal{A}^{\mathcal{P}}$ for objective ϕ_k^{a}. Therefore by Lemma 5, it has no winning strategy in A from $\gamma(s^{\mathsf{a}})$ neither for the objective ϕ_k. This means that $\gamma(s^{\mathsf{a}}) \in V_{k,-1}$.

$\boxed{\underline{V}_{k,0}}$ Note that by definition of the $\nu X.$ operator, $\underline{V}_{k,0} \subseteq F$. Thus, let us just show that $\gamma(F) \subseteq V_{k,0}$. Recall that $V_{k,0} = \mathsf{Win}_{\mathcal{P}\setminus\{k\}}(\mathsf{A}, \neg\phi_k) \cap \mathsf{Win}_{\mathcal{P}}(\mathsf{A}, \phi_k)$. Let $s \in \gamma(\underline{V}_{k,0})$. Then player k has no strategy in $\mathcal{A}^{\{k\}}$ for ϕ_k^{a}, hence, by Lemma 5, it cannot win A neither for ϕ_k from $\gamma(s)$. This shows that $\gamma(s) \subseteq \mathsf{Win}_{\mathcal{P}\setminus\{k\}}(\mathsf{A}, \neg\phi_k^{\mathsf{a}})$. Furthermore, the coalition $\emptyset$ has no strategy in $\mathcal{A}^{\mathcal{P}}$ for $\neg\phi_k^{\mathsf{a}}$, thus it does not have one neither in A for $\neg\phi_k$ from $\gamma(s)$. In other terms, $\gamma(s) \subseteq \mathsf{Win}_{\mathcal{P}}(\mathsf{A}, \phi_k)$.

□

We show that when playing according to $\underline{E}_k$, player k ensures staying in $\underline{V}$. This is proven in the following. Let us write $\gamma(\mathcal{E}) = \{(s, a) \mid (\alpha(s), a) \in \mathcal{E}\}$ for $\mathcal{E} \in \{\underline{E}_k, \overline{E}_k\}$.

Lemma 18. *For all games G, and players k, 1. $\gamma(\underline{E}_k \cap (\underline{V} \times \mathsf{Act})) \subseteq E_k \subseteq \gamma(\overline{E}_k)$. 2. For all $s^{\mathsf{a}} \in \mathsf{S}_k^{\mathsf{a}}$, there exist $a, a' \in \mathsf{Act}_k$ such that $(s^{\mathsf{a}}, a) \in \underline{E}_k$ and $(s^{\mathsf{a}}, a') \in \overline{E}_k$. 3. For all $(s^{\mathsf{a}}, a) \in \underline{E}_k$ with $s^{\mathsf{a}} \in \underline{V}$, we have $\mathsf{post}_{\Delta}(s^{\mathsf{a}}, a) \subseteq \underline{V}$.*

Proof. The inclusion $E_k \subseteq \gamma(\overline{E}_k)$ follows from the definition of $\overline{E}_k$, and by Lemma 6. It also follows that for all $s \in \mathsf{S}_k^{\mathsf{a}}$, there is $(s, a') \in \overline{E}_k$, since this is always the case for E_k.

Let (s^{a}, a) be an edge in $\underline{E}_k \cap (\underline{V} \times \mathsf{Act})$. Let s be a state in $\gamma(s^{\mathsf{a}})$. We have that $s \in \gamma(\underline{V}_{k,x})$ for some $x \in \{-1, 0, 1\}$ and by Lemma 6 $s \in V_{k,x}$. By definition of $\underline{E}_k$, for all t^{a} such that $\Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})$, $t^{\mathsf{a}} \in \underline{V}_{k,l}$ with $l \geq x$ and $s^{\mathsf{a}} \in \underline{V}_{k,x}$. By Lemma 6, we have that the value of all states in $\gamma(t^{\mathsf{a}})$ are at least as great as any state in $\gamma(s^{\mathsf{a}})$. By definition of Δ, $\alpha(\delta(s, a)) \subseteq \{t^{\mathsf{a}} \mid \Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})\}$. Therefore $\alpha(\delta(s, a)) \in \cup_{l \geq x} \underline{V}_{k,l}$, which means $\delta(s, a) \in \cup_{l \geq x} \gamma(\cup_{l \geq x} \underline{V}_{k,l}) \subseteq \cup_{l \geq x} V_{k,l}$ using Lemma 6. By definition of E_k this implies that $(s, a) \in E_k$.

It remains to prove that for all $s^{\mathsf{a}} \in \mathsf{S}_k^{\mathsf{a}}$, there is $(s^{\mathsf{a}}, a) \in \underline{E}_k$, and that if $s^{\mathsf{a}} \in \underline{V}$, then for all $(s^{\mathsf{a}}, a) \in \underline{E}_k$, $\Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})$ implies $t^{\mathsf{a}} \in \underline{V}$.

If $s^{\mathsf{a}} \in \mathsf{S}_k^{\mathsf{a}} \setminus \underline{V}$, then $(s^{\mathsf{a}}, a) \in \underline{E}_k$ for all $a \in \mathsf{Act}_k$ by definition. Let us now assume $s^{\mathsf{a}} \in \underline{V}$.

- If $s^{\mathsf{a}} \in \underline{V}_{k,-1}$, then By definition of $\underline{V}_{k,-1}$, we have that for all actions a, and all states t^{a}, if $\Delta^{\mathsf{a}}(s^{\mathsf{a}}, a, t^{\mathsf{a}})$ then $t^{\mathsf{a}} \in \underline{V}_{k,-1}$. Thus $(s^{\mathsf{a}}, a) \in \underline{E}_k$, and $t^{\mathsf{a}} \in \underline{V}_{k,-1}$ for any such t^{a}, so $t^{\mathsf{a}} \in \underline{V}$.
- If $s \in \underline{V}_{k,1}$, then there exists a such that $(s^{\mathsf{a}}, a, t^{\mathsf{a}}) \in \Delta^{\mathsf{a}}$ implies $t^{\mathsf{a}} \in \underline{V}_{k,1}$. So $(s^{\mathsf{a}}, a) \in \underline{E}_k$, and $t^{\mathsf{a}} \in \underline{V}_{k,1}$. Moreover this holds for all a with $(s^{\mathsf{a}}, a) \in \underline{E}_k$, since for such a, $(s^{\mathsf{a}}, a, t^{\mathsf{a}}) \in \Delta^{\mathsf{a}}$ implies $t^{\mathsf{a}} \in \underline{V}_{k,1}$ by definition of $\underline{E}_k$.
- If $s \in \underline{V}_{k,0}$, then by the greatest fixpoint defining $\underline{V}_{k,0}$, there exists $a \in \mathsf{Act}_k$ such that for all t^{a} with $\Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})$, $t^{\mathsf{a}} \in \underline{V}_{k,0}$. Conversely, for all $(s^{\mathsf{a}}, a) \in \underline{E}_k$, a ensures staying inside $\underline{V}_{k,0} \cup \underline{V}_{k,1}$. Thus for any such a, $(s^{\mathsf{a}}, a) \in \underline{E}_k$, and any t^{a}, $\Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})$ means $t^{\mathsf{a}} \in \underline{V}_{k,0}$.

□

Recall that $\underline{E}_k$ does not constrain the actions outside the set $\underline{V}$; thus strategies in $\mathsf{Strat}_k(\underline{E}_k)$ can actually choose dominated actions outside $\underline{V}$. To prove that $\mathsf{Strat}_k(\underline{E}_k)$ is an under-approximation of $\mathsf{Strat}_k(E_k)$ when started in $\underline{V}$, we need to formalize the fact that admissible strategies may choose arbitrary actions at states that are not reachable by any outcome. Intuitively, such strategies cannot be dominated since the dominated behavior is never observed.

For any strategy σ, let $\mathsf{Reach}(\mathsf{G}, \sigma)$ denote the set of states reachable from s_{init} by runs compatible with σ. We show that if one arbitrarily modifies an admissible strategy outside the set $\mathsf{Reach}(\mathsf{G}, \sigma)$, the resulting strategy is still admissible.

Lemma 19. *Let σ be a strategy in $\mathsf{Adm}_i(\mathsf{G})$ and σ' a strategy in $\Sigma_i(\mathsf{G})$. If for all histories h such that $\mathsf{last}(h) \in \mathsf{Reach}(\mathsf{G}, \sigma)$, $\Rightarrow \sigma(h) = \sigma'(h)$, then $\sigma' \in \mathsf{Adm}_i(\mathsf{G})$.*

Proof. For all profiles $\sigma_{-k} \in \Sigma_{-k}(\mathsf{G})$, we have $\mathsf{Out}_{\mathsf{G}}(\sigma_{-k}, \sigma) = \mathsf{Out}_{\mathsf{G}}(\sigma_{-k}, \sigma')$ so if σ' is dominated, then σ would also be dominated, which is a contradiction.

□

Lemma 7. *For all games G, and players k, $\mathsf{Strat}_k(E_k) \subseteq \gamma(\mathsf{Strat}_k(\overline{E}_k))$, and if $s_{\mathsf{init}} \in \gamma(\underline{V})$, then $\emptyset \neq \gamma(\mathsf{Strat}_k(\underline{E}_k)) \subseteq \mathsf{Strat}_k(E_k)$.*

Proof. Since $E_k \subseteq \gamma(\overline{E}_k)$ by Lemma 18, we have $\mathsf{Strat}_k(E_k) \subseteq \gamma(\mathsf{Strat}_k(\gamma(\overline{E}_k)))$.

Assume $s_{\mathsf{init}} \in \gamma(\underline{V})$. The fact that $\mathsf{Strat}_k(\underline{E}_k)$, thus also $\gamma(\mathsf{Strat}_k(\underline{E}_k))$ are non-empty follows from Lemma 18 too, since for any state s^{a} there is $a \in \mathsf{Act}_k$ with $(s^{\mathsf{a}}, a) \in \underline{E}_k$.

We prove that $\mathsf{Reach}(\mathcal{A}^{\mathcal{P}\setminus\{k\}}, \sigma) \subseteq \underline{V}$ for all $\sigma \in \mathsf{Strat}_k(\underline{E}_k)$. We already know, by Lemma 18, that for all $s^{\mathsf{a}} \in \underline{V}$, if $(s^{\mathsf{a}}, a) \in \underline{E}_k$ then all successors t^{a} with $\Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})$ satisfies $t^{\mathsf{a}} \in \underline{V}$. We are going to show that for all $s^{\mathsf{a}} \in \underline{V} \cap \mathsf{S}^{\mathsf{a}}_j$ with $j \neq k$, for all $a \in \mathsf{Act}_j$, $\Delta^{\mathsf{a}}(s^{\mathsf{a}}, a, t^{\mathsf{a}})$ implies $t^{\mathsf{a}} \in \underline{V}$.

Consider $s^{\mathsf{a}} \in \underline{V}$. If $s^{\mathsf{a}} \in \underline{V}_{k,1}$, then for all $a \in \mathsf{Act}$, $\Delta^{\mathsf{a}}(s^{\mathsf{a}}, a, t^{\mathsf{a}})$ implies that $t^{\mathsf{a}} \in \underline{V}_{k,1}$, since $\mathcal{P} \setminus \{k\}$ resolves non-determinism. The situation is similar if $s^{\mathsf{a}} \in \underline{V}_{k,-1}$; for all $a \in \mathsf{Act}_j$, $\Delta^{\mathsf{a}}(s^{\mathsf{a}}, a, t^{\mathsf{a}})$ implies $t^{\mathsf{a}} \in \underline{V}_{k,-1}$. If $s^{\mathsf{a}} \in \underline{V}_{k,0}$, then, by the definition of the outer fixpoint, for all $a \in \mathsf{Act}_j$, $\Delta^{\mathsf{a}}(s^{\mathsf{a}}, a, t^{\mathsf{a}})$ implies that $t^{\mathsf{a}} \in \underline{V}$.

Thus $\mathsf{Reach}(\mathcal{A}^{\mathcal{P}\setminus\{k\}}, \sigma) \subseteq \underline{V}$ for all $\sigma \in \mathsf{Strat}_k(\underline{E}_k)$. It then follows that $\mathsf{Reach}(\mathsf{G}, \gamma(\sigma)) \subseteq \gamma(\underline{V})$. So, by Lemma 19, and by the fact that $\gamma(\underline{E}_k) \subseteq E_k$, all strategies in $\gamma(\mathsf{Strat}_k(\underline{E}_k))$ are value preserving, which is to say, belong to $\mathsf{Strat}_k(E_k)$.

□

Help States

Lemma 20. *For all players k, $\gamma(\underline{H}_k) \subseteq H_k \subseteq \gamma(\overline{H}_k)$.*

Proof. Let $s^{\mathsf{a}} \in \underline{H}_k$, and let $a, b \in \mathsf{Act}$ two witnessing actions. For all $s \in \gamma(s^{\mathsf{a}})$, we have $\delta(s, a) \in \gamma(\mathsf{post}_{\Delta}(s^{\mathsf{a}}, a)) \subseteq V_{k,0} \cup V_{k,1}$ and $\delta(s, b) \in \gamma(\mathsf{post}_{\Delta}(s^{\mathsf{a}}, a)) \subseteq V_{k,0} \cup V_{k,1}$. Moreover $\alpha(\delta(s, a)) \in \mathsf{post}_{\Delta}(s^{\mathsf{a}}, a)$, $\alpha(\delta(s, b)) \in \mathsf{post}_{\Delta}(s^{\mathsf{a}}, b)$, and $\mathsf{post}_{\Delta}(s^{\mathsf{a}}, a) \cap \mathsf{post}_{\Delta}(s^{\mathsf{a}}, b) = \varnothing$, therefore $\alpha(\delta(s, a)) \neq \alpha(\delta(s, b))$ and thus $\delta(s, a) \neq \delta(s, b)$. Hence $s \in H_k$.

Now, consider any $s \in H_k$; and let $a, b \in \mathsf{Act}$ be such that $\delta(s,a), \delta(s,b) \in V_{k,0} \cup V_{k,1}$ and $\delta(s,a) \neq \delta(s,b)$. If we write $t^{\mathsf{a}} = \alpha(\delta(s,a))$ and $u^{\mathsf{a}} = \alpha(\delta(s,b))$, then $t^{\mathsf{a}}, u^{\mathsf{a}} \in \overline{V}_{k,0} \cup \overline{V}_{k,1}$, and $\Delta(s^{\mathsf{a}}, a, t^{\mathsf{a}})$, and $\Delta(s^{\mathsf{a}}, b, u^{\mathsf{a}})$; thus $\alpha(s) \in \overline{H}_k$. It follows that $H_k \subseteq \gamma(\overline{H}_k)$. □

Synthesizing AA-winning Strategies

Lemma 21. *We have* $\gamma(\underline{M'_k}) \subseteq M'_k \subseteq \gamma(\overline{M'_k})$.

Proof. We have $\gamma(\phi^{\mathsf{a}}_k) = \phi_k$ by assumption on γ. Thus, by Lemma 6,

$$\gamma((\mathtt{GF}(\overline{V}_{k,1}) \Rightarrow \phi^{\mathsf{a}}_k)) \subseteq \mathtt{GF}(V_{k,1}) \Rightarrow \phi_k) \subseteq \gamma((\mathtt{GF}(\underline{V}_{k,1}) \Rightarrow \phi^{\mathsf{a}}_k)).$$

Similarly, by Lemma 20, we get $\gamma\left(\mathtt{GF}(\overline{V}_{k,0}) \Rightarrow (\phi^{\mathsf{a}}_k \vee \mathtt{GF}(\underline{H}_k))\right) \subseteq \mathtt{GF}(V_{k,0} \Rightarrow (\phi_k \vee \mathtt{GF}(H_k)) \subseteq \gamma\left(\mathtt{GF}(\underline{V}_{k,0}) \Rightarrow (\phi^{\mathsf{a}}_k \vee \mathtt{GF}(\overline{H}_k))\right)$. It follows that $\gamma(\underline{M'}_k) \subseteq M'_k \subseteq \gamma(\overline{M'}_k)$.

□

The following lemma proves Theorem 5.

Lemma 22. *Let $k \in \mathcal{P}$ be a player and σ_k a strategy of player k. If $s^{\mathsf{a}}_{\mathsf{init}} \in \underline{V}$, and σ_k is winning for objective $\underline{\Omega'}_k$ in $\mathcal{A}'_k$, then $\gamma(\sigma_k)$ is winning for Ω'_k in G'_k.*

Proof. Let us rewrite

$$\underline{\Omega'}_i = \underline{M'}_i \wedge \left(\left(\mathtt{GF}(\mathsf{S}^{\mathsf{a}} \times \{0\}) \wedge (\bigwedge_{j \neq i} \overline{M'}_j \Rightarrow \phi^{\mathsf{a}}_i) \right) \vee \mathtt{GF}(\mathsf{S}^{\mathsf{a}} \times \{\top\}) \right).$$

Let σ_k be a winning strategy in $\mathcal{A}'_k$ for $\underline{\Omega'}_k$. We will show that $\mathsf{G}', \gamma(\sigma_k) \models \Omega'_k$.

Consider any run ρ of G'_k compatible with $\gamma(\sigma_k)$. By definition of $\gamma(\sigma_k)$, $\alpha(\rho)$ is a run of $\mathcal{A}'_k$ compatible with σ_k. Since σ_k is a winning strategy, $\alpha(\rho) \in \underline{M'}_k$, and by Lemma 21 $\rho \in M'_k$.

We now show that $\rho \in \mathtt{GF}(\mathsf{S} \times \{0, \top\})$. By assumption, we have $\mathcal{A}'_k, \sigma_k \models \mathtt{GF}(\mathsf{S}^{\mathsf{a}} \times \{0, \top\})$, which means that for all histories h^{a} of $\mathcal{A}'_k$ compatible with σ_k, $(\mathsf{last}(h^{\mathsf{a}}), \sigma(h^{\mathsf{a}})) \in \underline{E}_k$ (otherwise the transition relation of $\mathcal{A}'_k$ would lead to a $\bot$ state). Moreover, since $s^{\mathsf{a}}_{\mathsf{init}} \in \underline{V}$, it follows from Lemma 7 that $(\mathsf{last}(h), \gamma(\sigma)(h)) \in E_k$ for all histories h compatible with $\gamma(\sigma_k)$. Thus no state $(*, \bot)$ is reachable under $\gamma(\sigma)$ in G'_k.

Because of the structure of G'_k this means that ρ either visits states of $\mathsf{S} \times \{0\}$ or states of $\mathsf{S} \times \{\top\}$ infinitely often:

- If $\rho \in \mathtt{GF}(\mathsf{S} \times \{0\})$, then $\alpha(\rho) \in \mathtt{GF}(\mathsf{S}^{\mathsf{a}} \times \{0\})$; so $\alpha(\rho) \in \bigwedge_{j \neq k} \overline{M'}_j \Rightarrow \phi^{\mathsf{a}}_k$; it follows, by Lemma 21 and the compatibility of the abstraction with players' objectives, that $\rho \in \bigwedge_{j \neq k} M'_j \Rightarrow \phi_k$. Thus $\rho \in \Omega'_k$.
- Otherwise $\rho \in \mathtt{GF}(\mathsf{S} \times \{\top\})$, so $\rho \in \Omega'_k$.

Thus any outcome ρ of $\gamma(\sigma_k)$ belongs to Ω'_k which shows it is winning. □

E Algorithm for Assume-Guarantee Synthesis

The assume-guarantee-$\wedge$ rule was studied in [6] for particular games with three players. However, the given proofs are based on *secure equilibria* which do not actually capture assume-guarantee synthesis, so the correctness of the algorithm is not clear. Here, we give an alternative algorithm for deciding assume-guarantee-$\wedge$ for multiplayer games, and prove its correctness.

For any game G, and state s, we denote by G_s the game obtained making s the initial state. Assuming that for each player i has an objective ϕ_i which is prefix independent, let us define $W_i = \{s \in \mathsf{S} \mid \exists \sigma_i.\, G_s, \sigma_i \models \bigwedge_{j \in \mathcal{P} \setminus \{i\}} \phi_j \Rightarrow \phi_i\}$.

The following lemma gives a decidable characterization of assume-guarantee synthesis:

Lemma 23. *Let ϕ_i be a prefix-independent objective. Rule $\mathsf{AG}^\wedge$ has a solution if, and only if, there is a run ρ which visits only states of $\bigcap_{i\in\mathcal{P}} W_i$ and such that $\rho \models \bigwedge_{i\in\mathcal{P}} \phi_i$.*

Proof. $\boxed{\Rightarrow}$ Let $\sigma_\mathcal{P}$ be a solution of $\mathsf{AG}^\wedge$. Let ρ be its outcome. We have that $\rho \models \bigwedge_{i\in\mathcal{P}} \phi_i$ by hypothesis of $\mathsf{AG}^\wedge$. Let i be a player, we show that ρ only visits states of W_i. This is because σ_i is winning for $\bigwedge_{j\in\mathcal{P}\setminus\{i\}} \phi_j \Rightarrow \phi_i$. For all k, $\rho_{\leq k}$ is a finite outcome of σ_i, and the strategy played by σ_i after this history is winning for $\bigwedge_{j\in\mathcal{P}\setminus\{i\}} \phi_j \Rightarrow \phi_i$, which means that ρ_k belongs to W_i. Hence ρ satisfies the desired conditions.

$\boxed{\Leftarrow}$ If there is such a run ρ, we define the strategy profile $\sigma_\mathcal{P}$ to follow this run if no deviation has occurred and otherwise each player i plays a strategy which is winning for $\bigwedge_{j\in\mathcal{P}\setminus\{i\}} \phi_j \Rightarrow \phi_i$ if possible. We show that such a strategy profile satisfies the assumption of assume-guarantee. Obviously $\sigma_\mathcal{P} \models \bigwedge_{i\in\mathcal{P}} \phi_i$. Let ρ' be an outcome of σ_i and k the first index such that $\rho'_k \neq \rho_k$. The state $\rho'_{k-1} = \rho_{k-1}$ is not controlled by player i, because σ_i follows ρ. As ρ_{k-1} is in W_i and not controlled by player i, this means that $\rho'_k \in W_i$. Therefore σ_i plays a winning strategy from ρ'_k for the objective $\bigwedge_{j\in\mathcal{P}\setminus\{i\}} \phi_j \Rightarrow \phi_i$; thus ρ' satisfies this objective. Hence $\sigma_\mathcal{P}$ is a solution of $\mathsf{AG}^\wedge$.

□

We deduce a polynomial-space algorithm for the $\mathsf{AG}^\wedge$ rule with Muller objectives:

Theorem 8. *For multi-player games with Muller objectives, deciding whether $\mathsf{AG}^\wedge$ has a solution is* PSPACE*-complete.*

Proof. The algorithms proceed by computing the set W_i for each player i with an algorithm that computes winning regions and then checks whether there is an infinite run in the intersection $\bigcap_{i\in\mathcal{P}} W_i$ which satisfies $\bigwedge_{i\in\mathcal{P}} \phi_i$. This algorithm is correct thanks to Lemma 23.

This is in PSPACE because the objective $\bigwedge_{j\in\mathcal{P}\setminus\{i\}} \phi_j \Rightarrow \phi_i$ can be expressed by a Muller condition encoded by a circuit [20] of polynomial size. We can decide in polynomial space if a given state is winning for a Muller condition given by a circuit. Thus, the set $\bigcap_{i\in\mathcal{P}} W_i$ can be computed in polynomial space; let us denote by G' the game restricted to this set. The algorithm then consists in finding a run in G' satisfying $\bigwedge_{i\in\mathcal{P}} \phi_i$; that is, finding a run satisfying a Muller condition, which can be done in polynomial space.

□